



THE UNIVERSITY OF QUEENSLAND
AUSTRALIA

Physics of Low-Dimensional Ultracold Bose Gases

Christopher James Foster

BSc(Hons), BMath

*A thesis submitted for the degree of Doctor of Philosophy at
The University of Queensland in 2011*

School of Mathematics and Physics

Abstract

In this thesis we investigate the properties of ultracold Bose gases in one and two dimensions. Low-dimensional systems have several striking differences from their three-dimensional counterparts, most notably the absence of Bose-Einstein condensation (BEC) in the infinite homogeneous case. Recent experimental progress has brought low-dimensional systems within reach in the laboratory, and we provide numerical simulations with experimentally relevant parameters. We present simulations of the Berezinskii-Kosterlitz-Thouless (BKT) phase transition in two dimensions and investigate equations of motion for a Bose gas constrained to one dimension.

Recent experiments have attributed superfluidity in two-dimensional systems to the BKT phase transition. We perform classical field simulations using the projected Gross-Pitaevskii equation (PGPE) formalism to model the two-dimensional Bose gas at finite temperature. We confirm the presence of the BKT phase via the observation of two unique features: algebraic decay of the first-order correlation function; and vortex pair unbinding. Unbinding of vortex pairs at the BKT transition is clearly demonstrated via a coarse-graining procedure which reveals unpaired vortices. The BKT transition temperature identified via correlations and vortex behaviour agrees well with the temperature deduced from a calculation of the superfluid fraction. Surprisingly, we observed no separation between the temperature of the BEC transition — which is present due to the finite simulation size — and the BKT transition. We relate our results to experimental observations and show that an interpretation based on BKT physics is justified.

In investigating the two-dimensional system we found it necessary to compute the superfluid fraction. Calculating the superfluid fraction is a delicate procedure because it involves connecting the macroscopic quantum phenomenon of superfluidity to a detailed microscopic simulation. We present an efficient method which overcomes this difficulty using a tensor decomposition of the momentum density autocorrelations. Our method gives results consistent with the other physical properties of the BKT phase in two-dimensional systems.

One-dimensional effective equations allow for efficient simulation of very elongated systems of ultracold Bose gases. We use a Gaussian ansatz and Lagrangian approach to derive an effective equation for a Bose gas constrained to one dimension at zero temperature. In some respects our method outperforms alternative one-dimensional equations such as the non-polynomial Schrödinger equation. However, our scheme is found to be inherently unstable in

a wide range of cases. We analyse the instability and find that it is an inherent feature of a whole class of nonlinear ansätze.

Declaration by author

This thesis is composed of my original work, and contains no material previously published or written by another person except where due reference has been made in the text. I have clearly stated the contribution by others to jointly-authored works that I have included in my thesis.

I have clearly stated the contribution of others to my thesis as a whole, including statistical assistance, survey design, data analysis, significant technical procedures, professional editorial advice, and any other original research work used or reported in my thesis. The content of my thesis is the result of work I have carried out since the commencement of my research higher degree candidature and does not include a substantial part of work that has been submitted to qualify for the award of any other degree or diploma in any university or other tertiary institution. I have clearly stated which parts of my thesis, if any, have been submitted to qualify for another award.

I acknowledge that an electronic copy of my thesis must be lodged with the University Library and, subject to the General Award Rules of The University of Queensland, immediately made available for research and study in accordance with the *Copyright Act 1968*.

I acknowledge that copyright of all material contained in my thesis resides with the copyright holder(s) of that material. Where appropriate I have obtained copyright permission from the copyright holder to reproduce material in this thesis.

Publications during candidature

Peer reviewed papers:

- [46] C. J. Foster, P. B. Blakie, and M. J. Davis. Vortex pairing in two-dimensional Bose gases. *Phys. Rev. A*, 81:023623, 2010.
- [18] E. G. Cavalcanti, C. J. Foster, M. D. Reid, and P. D. Drummond. Bell inequalities for continuous-variable correlations. *Phys. Rev. Lett.*, 99:210405, 2007.
- [84] P. E. M. F. Mendonça, R. d. J. Napolitano, M. A. Marchioli, C. J. Foster, and Y.-C. Liang. Alternative fidelity measure between quantum states. *Phys. Rev. A*, 78:052330, 2008.

Publications included in this thesis

Ref. [46] is incorporated as chapters 3, 4 and appendix B of the thesis. I was responsible for writing the majority of this paper and performed all technical calculations. P. Blair Blakie assisted in the writing of the introductory and concluding parts of the paper. Matthew J. Davis guided the research and helped with interpretation of results.

Ref. [18] is reproduced verbatim in appendix D. Eric G. Cavalcanti was responsible for writing most of the paper and performing most of the technical calculations. I contributed the no-go proof on pages three and four in collaboration with Eric, and helped with the final stages of writing. Margaret D. Reid analysed the effect of detector inefficiencies. Peter D. Drummond conceived and guided the research.

Ref. [84] is also reproduced verbatim in appendix D. Paulo E. M. F. Mendonça wrote the initial draft of the manuscript containing several of the presented scientific results. Reginaldo d. J. Napolitano obtained the results of section III and the proof of proposition III.2. Marcelo A. Marchioli contributed with discussions and improvements of the text, and collaborated with Paulo in conjecturing inequality (29) of the paper. I contributed the numerical results and written description of section IV; I also assisted in revising the manuscript as a whole. Yeong-Cherng Liang conceived section IV including initial Matlab code, and contributed to the revision of the manuscript as a whole.

Contributions by others to the thesis

Simon A. Haine provided initial ideas and Matthew J. Davis provided initial ideas and general guidance for the work described in chapter 5. Vivien J. Challis, S. A. Haine and M. J. Davis proofread the thesis and provided suggestions.

Statement of parts of the thesis submitted to qualify for the award of another degree

The two works incorporated verbatim in appendix D were previously incorporated into PhD theses by their respective primary authors: Ref. [18] was submitted by fellow author Eric G. Cavalcanti as part of his PhD at the University of Queensland, degree awarded 19 May 2008; Ref. [84] was submitted by fellow author Paulo E. M. F. Mendonça as part of his PhD at the University of Queensland, degree awarded 11 August 2009.

Acknowledgements

There are many people who contributed positively to this thesis in one way or another. The most obvious of course is my supervisor Matthew Davis. Matt has been involved with the project from the beginning, and I am deeply grateful for his unwavering support and good humoured patience.

Thanks are also due to Blair Blakie, Ashton Bradley, Joel Corney, Peter Drummond, Simon Haine and Tod Wright, all of whom were involved in my supervision at some stage, either officially or unofficially. I am particularly grateful to Blair for providing the final impetus to publish the work of chapter 3, and to Simon for finding the time to proofread much of this thesis and provide thoughtful suggestions.

My PhD experience would have been vastly inferior without the friendship and camaraderie of my fellow students. Most obviously there were “the usual suspects”: my office mates Eric Cavalcanti, Andy Ferris, Geoff Lee, Yeong-Cherng Liang, Terry McRae and Paulo Mendonça, who together made the office a lively and intellectually stimulating environment. Memories which will not be easily forgotten include the animated discussions with Eric on the nature of consciousness, coding and juggling adventures with Andy, and many fascinating fragments of the home cultures of Eric, Paulo, and Yeong-Cherng. I also enjoyed interacting with other students within and outside the department, including Andrew Sykes, Tim Vaughan, Rodney Polkinghorne, Mark de Burgh and many more. Most of these friends have already moved on from UQ, so I can say with certainty that they are sorely missed.

The research group at the UQ node of the Australian Research Council Centre of Excellence for Quantum-Atom Optics (ACQAO) also deserves my thanks for hosting me during my PhD. The members of the group have always provided a good environment to learn about theoretical physics, but also about the *process* of doing theoretical physics. In addition, ACQAO as an organisation provided funding for both part of my scholarship and travel to various conferences. I am particularly grateful for the opportunity to travel to conferences in France and Germany during early 2005.

Finally, I thank my family and particularly my wife Vivien for emotional and practical support over these last several years. To Vivien I would like to say — “Thank you for your love, thank you for your support, and thank you for being my last line of defence against sloppy definitions and awkward sentences!” It is quite possible that this thesis would not exist without Vivien, and with that in mind, I dedicate it to her.

Keywords

Bose-Einstein condensation, ultracold Bose gases, classical field methods, superfluidity, low-dimensional systems, Berezinskii-Kosterlitz-Thouless phase transition, projected Gross-Pitaevskii equation

Australian and New Zealand Standard Research Classifications (ANZSRC)

ANZSRC code: 020601, Degenerate Quantum Gases and Atom Optics, 100%

Fields of Research (FoR) Classification

FoR code: 0206, Quantum Physics, 100%

Contents

List of Figures	xiii
List of Abbreviations	xiv
1 Introduction	1
1.1 Bose statistics and condensation	4
1.2 Low-dimensional systems	6
1.3 Outline of the thesis	9
2 Background Theory	11
2.1 Bose statistics and BEC in noninteracting systems	11
2.1.1 The Bose distribution	11
2.1.2 Condensation in the 3D homogeneous Bose gas	13
2.2 Field theory and interacting Bose gases	16
2.2.1 Second quantisation	16
2.2.2 The low-energy Hamiltonian	21
2.2.3 Operator equations of motion	22
2.3 The Gross-Pitaevskii equation	23
2.3.1 The Dirac-Frenkel time dependent variational principle	24
2.3.2 Deriving the GPE	25
2.3.3 Validity of the GPE	27
2.4 The Projected Gross-Pitaevskii equation	29
2.4.1 Conceptual introduction	29
2.4.2 Derivation of the PGPE	30
2.4.3 Ergodicity and thermal averages	34
2.4.4 Treatment of the I region	36
2.4.5 Comparison with the GPE	37

3	Vortex pairing in two-dimensional Bose gases	39
3.1	Introduction	39
3.2	Formalism	40
3.2.1	Review of BKT physics	40
3.3	Method	42
3.3.1	c-field and incoherent regions	42
3.3.2	Mean-field treatment of I region	45
3.3.3	Equilibrium configurations with fixed T and N	45
3.4	Results	46
3.4.1	Simulation of expanded interference patterns between two systems	46
3.4.2	Condensate and superfluid fractions	47
3.4.3	First-order correlations — algebraic decay	49
3.4.4	Vortices and pairing	53
3.5	Conclusion	60
4	Superfluid fraction and the PGPE	61
4.1	Introduction	61
4.2	Momentum density correlations and the superfluid fraction	64
4.3	Numerical procedure	68
4.4	Discussion	69
5	Effective 1D equations for trapped Bose-Einstein condensates	71
5.1	Introduction	71
5.1.1	Previous work	72
5.2	Formalism	75
5.2.1	The Ansatz	75
5.2.2	Equations of motion	75
5.2.3	Conservation of normalisation and conserved current	77
5.2.4	Ground states	78
5.3	Numerical Simulations	79
5.3.1	Ground states	79
5.3.2	Test case: soliton formation	80
5.3.3	Transverse shock wave formation	83
5.4	Conclusion	85
6	Conclusion	87

References	88
A Mathematical techniques for classical field theory	99
A.1 The Wirtinger calculus: derivatives of nonholomorphic functions	99
A.1.1 Properties of the Wirtinger derivative	101
A.1.2 Computing with Wirtinger derivatives	103
A.2 Functional derivatives	104
A.2.1 An example	104
A.2.2 General definition	106
A.2.3 Calculating with functional derivatives	106
B Details of two-dimensional simulations	109
B.1 Simulation using the PGPE	109
B.1.1 Hartree-Fock-Bogoliubov analysis	109
B.1.2 Initial conditions for fixed total number	110
B.1.3 Initial conditions for given $E_{\mathbf{C}}$ and $N_{\mathbf{C}}$	111
B.2 \mathbf{I} region integrals	112
B.3 Vortex detection	113
C An imaginary time method with fixed normalisation	115
C.1 The usual imaginary time method	115
C.2 The imaginary time method with varying μ	117
D Additional work	119

List of Figures

1.1	Bose statistics versus distinguishable statistics in a two mode system	5
3.1	Simulated interference fringes showing zipper patterns due to vortices	48
3.2	Condensate and superfluid fractions as a function of temperature	49
3.3	Detail of the superfluid fraction near the transition temperature	50
3.4	Power law fits of $g^{(1)}$ in two dimensions	51
3.5	Two methods for determining the algebraic decay coefficient for $g^{(1)}$	52
3.6	Phase profile of a c-field showing vortex pairing at two temperatures	55
3.7	Number of paired and unpaired vortices as a function of temperature	56
3.8	Vortex pair distribution function at various temperatures	57
3.9	Coarse-graining procedure for detecting unpaired vortices	58
3.10	Comparison of vortex unpairing measures	59
4.1	Thought experiment used in deriving the superfluid density	65
4.2	Extrapolation of transverse and longitudinal momentum correlations	69
5.1	Solitons on a constant density background	81
5.2	Comparison of simulation methods for the soliton waveguide problem	82
5.3	Discontinuities in the transverse Gaussian width	84
5.4	Radial structure of 3D waveguide simulation	85
5.5	Fits to the 3D wavefunction on both sides of a width discontinuity	86

Introduction

In this thesis we investigate the behaviour of dilute, degenerate Bose gases confined to one and two dimensions. This introductory chapter is mostly dedicated to explaining the physics contained in the previous sentence: The quantum statistics of bosons, the nature of degeneracy, and the physical differences between three and lower dimensions.

It is a surprising fact that elementary particles of the same type are not distinguishable, even in principle. The first hints of the importance of this fact appeared even before the advent of quantum mechanics as a resolution to the Gibbs paradox of classical statistical mechanics: The calculated entropy of an ideal gas is not extensive¹ unless the atoms are treated as indistinguishable. As it turns out, there are exactly two consistent statistical behaviours for indistinguishable quantum particles in three dimensions². Bose discovered the first of these in 1924 [17] in a successful attempt to derive Planck’s blackbody radiation formula from first principles. Bose had invented a new way to count multi-photon states, which Einstein immediately applied to atoms to predict the equation of state [36], and in 1925 the “condensation” [37] of what became known as the ideal Bose gas. Bose-Einstein condensation (BEC) was invoked in the following decades to explain various condensed matter phenomena including superfluidity in liquid helium and low temperature superconductivity. In 1926 Fermi [40] and Dirac [33] discovered the second type of quantum statistics, motivated partly by the behaviour of electrons as newly articulated in the Pauli exclusion principle. Particles obeying Bose-Einstein and Fermi-Dirac statistics are now known as *bosons* and *fermions* respectively.

We reflect briefly on what exactly we mean by “particle statistics”. Statistical mechanics starts from a probability distribution over the system microstates³ and uses this distribution to calculate expected properties of the system. The distribution is specified by some macroscopic

¹A measured quantity is *extensive* if it is directly proportional to the size of the system; for a new system composed of two exact copies of the system joined together the extensive quantity is doubled.

²In two dimensions there is in fact a continuum of possibilities known as anyonic statistics. The systems considered in this thesis are embedded in real three-dimensional space however, so anyonic statistics are irrelevant at the level of the constituent particles.

³A microstate is a full microscopic description of the details of a physical system; a state of maximal information.

variables (for example the temperature) but before the distribution can be given, the possible microstates must be described. Bosons, fermions and classical particles differ in exactly which microstates are possible. For a noninteracting system such as the ideal gas, the state of the full system may be given by specifying which single-particle states, or *modes*⁴, are occupied. Here is where the difference arises — any number of bosons may occupy a given mode, while the occupation is constrained to zero or one particle for fermions. Labelling the states of the full system using only the occupation of modes incorporates particle indistinguishability. For classical distinguishable particles, we would also need to keep track of *which* particles are in each mode, in addition to how many.

Quantum statistical behaviour is connected to intrinsic angular momentum by the celebrated spin statistics theorem: particles of half-integer spin (in multiples of \hbar) are fermions while particles with integer spin are bosons. Although the elementary constituents of normal matter are all fermions, bosonic statistics may be observed in composite particles — atoms for example — composed of even numbers of fermions. For this to be true, it is sufficient that the interactions between composites be at much lower energy than the internal energy levels [35].

In order to observe the consequences of quantum statistics a system must be *degenerate*. In this context “degenerate” means that the number of energetically accessible modes is similar to or smaller than the total number of particles; it is only with this kind of crowding of the available states that the differences in counting microstates become apparent. Degeneracy may be reached by either increasing the density, which increases the number of particles per state, or reducing the temperature which reduces the number of accessible states. For the particular case of an ideal gas, it is convenient to discuss degeneracy in terms of the de Broglie wavelength: At a given temperature T , the expected energy of a classical ideal gas particle is $\frac{1}{2}k_B T$ for each degree of freedom, according to the equipartition theorem, where k_B is Boltzmann’s constant. In three dimensions we therefore have $\langle p^2/2m \rangle = \frac{3}{2}k_B T$, and using de Broglie’s formula $\lambda = h/p$ gives a value $\lambda = h/\sqrt{3mk_B T}$ for the typical quantum wavelength. It is conventional⁵ to define the *thermal de Broglie wavelength* as

$$\lambda_{\text{dB}} = \sqrt{\frac{2\pi\hbar^2}{mk_B T}}. \quad (1.1)$$

With this definition the number of accessible states can be estimated as $Z = V/\lambda_{\text{dB}}^3$ where V is

⁴It is easy to confuse the terminology when referring to “single-particle states” versus “states of the system as a whole”. We hope to avoid this problem by following a common practice from optics where the single-particle states are referred to as *modes*.

⁵The convention for λ_{dB} is selected to absorb the constants in the expression for the ideal gas partition function.

the volume of the system⁶. Quantum degeneracy occurs when the number of particles is greater than or approximately equal to the number of accessible states, that is, $N \gtrsim Z$. Rearranging, we arrive at the degeneracy condition

$$n \gtrsim \lambda_{\text{dB}}^{-3} \quad (1.2)$$

for the density $n = N/V$, which has a nice physical interpretation: Degeneracy occurs when the size of the quantum wavepacket is comparable to or larger than the interparticle spacing.

Apart from Bose's work, the first historical uses of quantum statistics were to compute properties of the ideal gas in the quantum degenerate regime. Nevertheless, there was no experimentally accessible system with both weak interactions and strong quantum statistical behaviour until 1995 when the first degenerate, dilute Bose gases were created in the laboratory [4, 25]. This feat was followed shortly afterwards with the creation of a degenerate Fermi gas in 1999 [32]. There were several good reasons why it took seventy years for experimental systems to catch up with the theory. On the one hand, quantum degeneracy occurs in many sorts of condensed matter systems, including striking examples such as superfluidity in liquid ^4He . However, these systems have strong interactions between particles which renders the simplest theories useless from a quantitative perspective. On the other hand, gases with weak interactions are nowhere near degeneracy at easily accessible temperatures, and tend to form liquids and solids when cooled. At low densities the formation of liquids or solids from atomic gases is suppressed, and cooling to extremely low temperatures — of the order of 100 nK — allows experiments to reach the degenerate regime. Perfecting the experimental tools necessary to trap and cool atomic clouds to such extraordinarily low temperatures was a major undertaking.

From now on we focus exclusively on bosons which are the topic of this thesis. The typical ultracold boson experiment begins with a gas of neutral atoms — ^{87}Rb for example — in a vacuum chamber. In the first stage the atoms are captured in a magneto-optical trap and cooled via laser cooling to temperatures of order 100 μK [98]. Atoms are then loaded into a tighter magnetic or optical dipole trap and further cooled by successively removing the most energetic atoms — a process known as evaporative cooling. During this process the gas undergoes condensation to form a BEC at temperatures of order 1 μK to 100 nK, with total number of atoms ranging from 10^8 to 10^3 [96, §1.1]. The trapping potential is usually well approximated by a parabola, with the spatial extent of the atomic cloud on the order of 10–100 μm , depending on the trap anisotropy [5].

The precision and control available to experiments has continued to improve over the last 15 years. Experiments can now control the interaction strength via Feshbach resonances, and

⁶Here Z is actually the canonical partition function; if the energy scale is adjusted such that the lowest energy state has energy zero, Z counts the accessible states. (See, for example, [115, §6.1].)

detailed control of the trapping potential is available using combinations of lasers and magnetic fields. Ultimately, the experimental accessibility of weakly interacting and precisely controllable quantum gases has been a great resource for basic quantum physics: It has allowed rigorous comparison and evaluation of first principles approaches to quantum field theory.

1.1 Bose statistics and condensation

The phenomenon of BEC is a phase transition, with a normal gaseous phase at higher temperature, and a fluid with macroscopic quantum behaviour at low temperature. In the low temperature phase, a macroscopic number of particles all occupy the same quantum state, in a sense to be made more precise below. For the particular case of the noninteracting homogeneous gas in three dimensions, the expected occupation of the ground state is

$$N_0 = N \left[1 - \left(\frac{T}{T_c} \right)^{3/2} \right], \quad (1.3)$$

where T_c is the transition temperature. We discuss the value of T_c and provide a derivation of this basic relation in section 2.1.2.

Unlike most phase transitions, BEC occurs even in systems without interparticle interactions: it is a purely statistical effect. To understand qualitatively how this happens, we consider the effect of Bose statistics on a two mode toy system containing exactly N particles. For definiteness, let us think of these modes as a pair of “left” and “right” potential wells.

First let us consider the case where each well has the same energy, and count the number of states available to the system as a whole. For classical distinguishable particles, each particle may be in the left or right well independently of the rest, so adding a particle to the system multiplies the number of available states by two; the total number of states for N particles is then 2^N . On the other hand, for indistinguishable particles a state of the system as a whole is fully specified by listing the number of particles in each well. The number of particles in the left well, N_L , is between 0 and N particles, and N_R is always equal to $N - N_L$, so there are only $N + 1$ states.

If both modes have the same energy then all states of the system have the same energy and are therefore equally likely⁷. The probability of having all N particles in the left well is then $1/2^N$ for the distinguishable case, but $1/(N + 1)$ for indistinguishable particles. From this simple example, we can already see how Bose statistics exponentially enhances the probability of finding the system with all particles in a single well, simply because of the way states are

⁷The assumption that states with equal energy are equally probable is the founding assumption of statistical mechanics, see, for example, Ref. [115, §7.1].

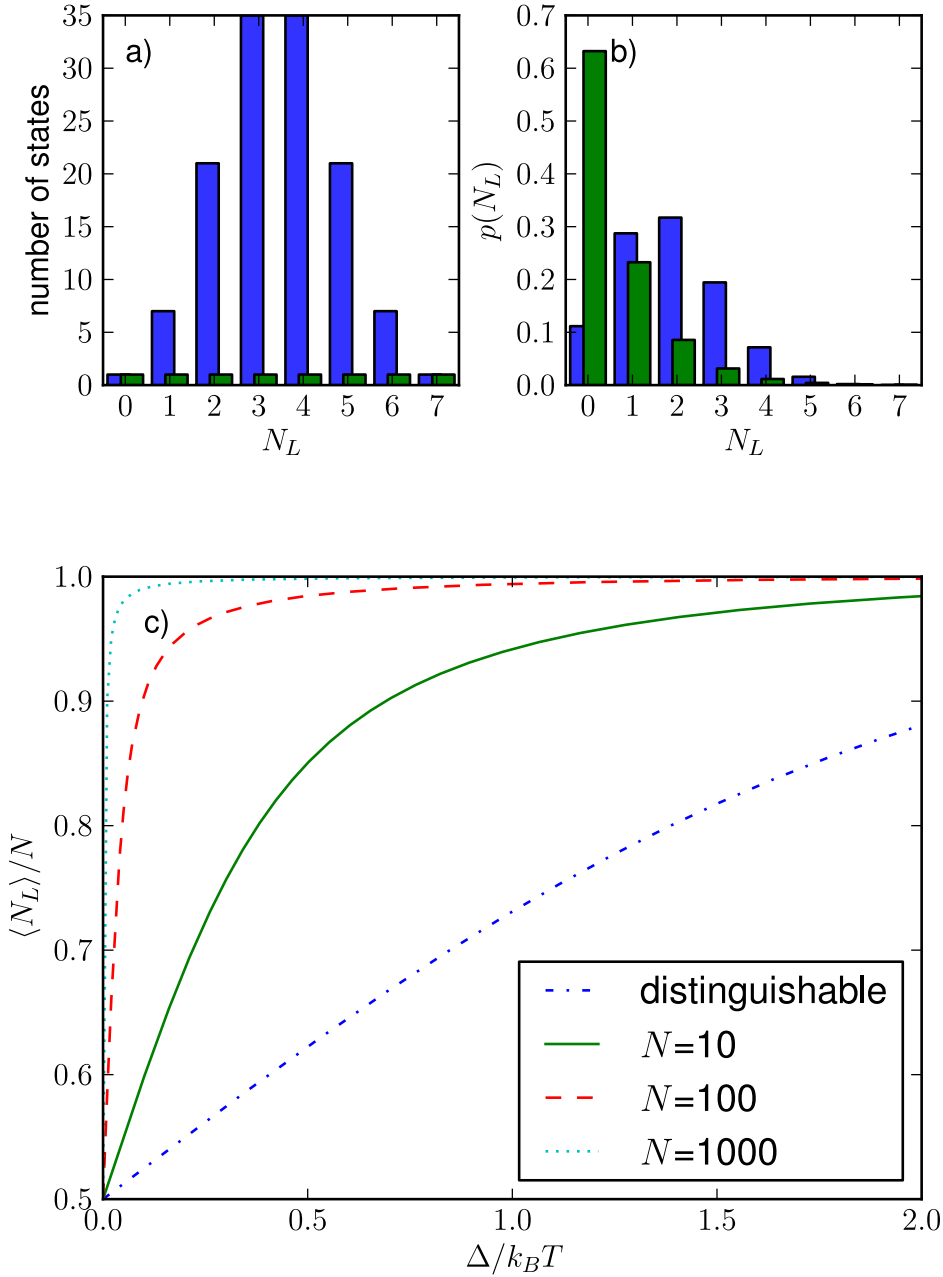


Figure 1.1: Bose statistics versus classical distinguishable statistics in a two mode toy system. (a) Count of system states with a given number of particles in the left well for the distinguishable case (blue) and bosonic case (green). This is proportional to the probability density $p(N_L)$ when $\Delta = 0$. (b) Probability density $p(N_L)$ for finding the system with N_L particles in the left well, when $\Delta/k_B T = 1$ [colours as in (a)]. Both (a) and (b) show statistics for a system of $N = 7$ atoms. (c) Expected fraction of particles in the left well as a function of $\Delta/k_B T$ for the bosonic case with three different atom numbers. The distinguishable case is included for reference and is the same for *any* N .

counted. One way to visualise this is to plot the number of states with a particular N_L as in Fig. 1.1(a) — even for the very small value of $N = 7$ shown, there are many more distinguishable states where the particles are evenly distributed between the wells.

If we now let the right well have energy Δ greater than the left, the probability distribution $p(N_L)$ is skewed proportionally to $e^{-N_L\Delta/k_B T}$, as shown in Fig. 1.1(b)⁸. Nevertheless, the distinguishable case may still have a peak at nonzero N_L due to the statistical pressure toward equal numbers; this effect becomes much more pronounced for larger, more realistic values of N . Taking the expected value of N_L from the probability distribution, we can also calculate⁹ a “condensate fraction” — the expected fraction of atoms in the lowest energy state, $\langle N_L \rangle / N$. Figure 1.1(c) shows the condensate fraction $\langle N_L \rangle / N$, as a function of $\Delta/k_B T$, for various total number N . It is clear that for bosons the condensate fraction rapidly converges to 1 as N increases, regardless of the size of the energy gap between the wells. On the other hand, for distinguishable particles the fraction in the left well is independent of N . It is worth noting that in this toy system — or any system with a finite number of modes — there must always be a mode with an extensive population, so strictly speaking the existence of condensation is a trivial fact in this case. Even so, it is instructive to observe the differences between the bosonic and distinguishable cases.

An infinite homogeneous gas has an infinite number of modes, so the existence of a nonzero condensate fraction in the lowest energy mode is a nontrivial fact in this case (see section 2.1.2). The existence of BEC in a system with infinitely many modes turns out to depend on the density of states at low energies. In the three-dimensional homogeneous case, there are “sufficiently few” low-energy modes meaning that particles are forced to gather in the lowest energy mode and BEC occurs. However, in one and two dimensions the density of low-energy modes is larger and BEC is forbidden, at least in the homogeneous case. We emphasise that the addition of a trapping potential modifies the density of states, which can lead to BEC even in low dimensions.

1.2 Low-dimensional systems

The physical behaviour of one- and two-dimensional systems can be qualitatively different from a three-dimensional (3D) system containing the same constituent particles. A dramatic example is the homogeneous ideal Bose gas which undergoes the BEC phase transition in 3D, but never condenses in lower dimensions at any nonzero temperature. This absence of BEC is a specific

⁸The Boltzmann factor $e^{-N_L\Delta/k_B T}$ arises from analysing the system in the canonical ensemble, see, for example, Ref. [115, Ch. 6].

⁹The calculation of $\langle N_L \rangle / N$ is immediate in the canonical ensemble via direct summation over states.

case of the more general Mermin-Wagner-Hohenberg theorem [86, 66], which forbids long-range order in a whole class of one- and two-dimensional systems.

The qualitative differences between three and lower dimensions are due to both geometric and topological effects. For example, the existence of BEC depends critically on the density of states at low energy, which for the homogeneous case scales as the surface area of a $(d-1)$ -sphere in momentum space, where d is the dimension; this is a geometric effect (see section 2.1.2). On the other hand, the nature and existence of elementary topological excitations such as vortices also depends on the dimensionality [19, Ch. 9] and these excitations have a significant effect on the thermodynamics of the system.

At first sight one might question the relevance of low-dimensional systems, given that the world is really three-dimensional. However, many 3D systems can be made to act as if they are low-dimensional by tight confinement. When the confinement is sufficiently strong and the temperature low enough, the transverse degrees of freedom are effectively “frozen out” and all dynamics occurs within the plane or line. More precisely, when interactions are unimportant it is clear that freeze out occurs for $\hbar\omega \gg k_B T$, where we have compared one quantum of transverse excitation energy $\hbar\omega$ with the characteristic magnitude $k_B T$ of the thermal fluctuations. This estimate holds even when interactions are important, as discussed in Ref. [96, Ch. 15].

Ref. [54] presents a useful characterisation of low-dimensional regimes in terms of length scales: Let us consider a BEC in a cylindrically symmetric system, where the extent of the trapped cloud is given by radii R_\perp and R_z in the radial and longitudinal directions respectively. As discussed in Ref. [54], there are two additional length scales relevant in this problem, the healing length¹⁰ ξ and the scattering length a_s . In many experiments, we have the situation $R_\perp, R_z \gg \xi \gg a_s$, in which case the dynamics are fully three-dimensional. In contrast, the 2D and 1D cases correspond to the regimes $R_\perp \gg \xi > R_z$ and $R_z \gg \xi > R_\perp$, where the BEC is conventionally called “pancake-shaped” and “cigar-shaped” respectively. In all these cases we assume that $R_\perp, R_z \gg a_s$ so that the scattering remains three-dimensional, even when the dynamics is two- or one-dimensional. This is called the *quasi*-low-dimensional regime [54].

We briefly discuss some unique properties of two-dimensional systems. As noted above, condensation does not happen in a homogeneous 2D system. However, the 2D Bose gas supports topological defects in the form of vortices, and in the presence of interactions can instead undergo a Berezinskii-Kosterlitz-Thouless (BKT) [9, 75, 104] transition to a quasi-coherent superfluid state. By quasi-coherent we mean a state with algebraic¹¹ decay of spatial correlations,

¹⁰The healing length $\xi = 1/\sqrt{8\pi n a_s}$ is the length scale on which a condensate recovers (“heals”) from the effects of a local perturbation, where n is the density. ξ may be estimated as the length scale on which the kinetic term in the Gross-Pitaevskii equation (described in section 2.3) is equal to the size of the interaction term far from the perturbation. For details see Ref. [96, §6.4].

¹¹“Algebraic” decay is the name conventionally used in this context for a simple power law decay.

rather than the traditional off diagonal long-range order of a coherent condensate [124], or the short range exponential decay characteristic of a high temperature phase. This is one of several unique features of the BKT phase.

The most striking feature of the BKT transition is the pairing of topological defects in the low temperature phase. In the superfluid case this corresponds to pairing of quantised vortices of opposite rotation, which begin to unbind as the temperature is increased through the transition; the high temperature phase is a plasma of bound and unbound vortices [75].

A simple argument shows why vortex unbinding is related to the destruction of superfluidity above the BKT transition. Consider a channel filled with a superfluid, and transport a single unpaired vortex from one boundary of the channel to the other. This imparts a single quantum of velocity to the fluid within the channel¹², and provides a mechanism for the decay of supercurrents. On the other hand, transport of *pairs* of vortices does not impart any net momentum to the channel. We defer further presentation of the aspects of the BKT transition relevant to 2D Bose gas experiments until chapter 3.

In addition to 2D BKT physics, we also consider BECs in 1D systems in this thesis. In contrast to the 2D case, we make no particular attempt to investigate features unique to one dimension, since our focus is on developing a numerical approximation technique. However, we do rely on one uniquely 1D feature for testing our numerical method: In 1D there are stable soliton solutions with well understood behaviour; in higher dimensions such solitons would be unstable, and eventually decay via the “snake instability” [3].

Experiments in cold gases have been able to access the low-dimensional regime since 2001 when Görlitz *et al.* created both 1D and 2D systems using magnetic and optical traps respectively [54]. The BKT transition was first observed in liquid helium thin films [11], but observing aspects of BKT physics in an ultracold Bose gas proved difficult until a series of experiments by the Dalibard group at ENS in Paris [121, 60, 77]. In the Paris experiments, an optical lattice was used to create two parallel pancake-shaped BECs. Using these as mutual phase references allowed unpaired vortices to be detected in an interference experiment [121]. Various aspects of BKT physics including the algebraic decay of correlations were subsequently observed in Refs. [60, 77]. Further experimental work by other groups has followed [117, 23], further adding to the evidence of BKT physics in the 2D cold gas system.

¹²It is easiest to see how the single quantum is imparted by considering a toroidal channel where the superflow represents angular momentum, and transporting a vortex from the centre to the outside.

1.3 Outline of the thesis

After the current introductory chapter, some background theory is introduced in chapter 2, including the two main pieces of formalism used later in the thesis: The time dependent Gross-Pitaevskii equation (to be used in chapter 5) is introduced in section 2.3; the classical field methods for finite temperature calculations (in particular the projected Gross-Pitaevskii equation, to be used in chapter 3) are introduced in section 2.4. An attempt has been made to keep the background self-contained, starting from the level of an advanced undergraduate. For this reason, a section on basic Bose statistics is included, along with a discussion of how nonrelativistic quantum field theory arises as a generalisation of first-quantised many-body quantum mechanics.

The two-dimensional experiments discussed in the previous section present a challenge for theory. One particular issue is to understand the relation between the BKT phase and BEC, which exists in real experiments due to the harmonic trapping potential. In chapter 3 we make use of classical field methods to simulate the finite temperature physics of 2D systems, with a view to understanding this question and to understand which aspects of BKT physics may be observed in current experiments.

As part of this work, we developed a method for calculating the superfluid fraction from a classical field simulation. This applicability of this method is more general than the specific use to which it is put in chapter 3, so we describe the derivation in isolation as chapter 4.

We change focus in chapter 5 to examine effective equations for 1D systems. Motivated by the desire to study dispersive shock waves in high resolution and perform expansions of suddenly untrapped atomic clouds, we consider a variational ansatz for the quasi-1D time evolution. We implement the method numerically and characterise the behaviour using several test problems.

The thesis concludes with a summary in chapter 6. Several appendices are included: appendix A describes subtleties surrounding differentiation of Lagrangian functionals with respect to complex fields. Appendices B and C describe additional details pertaining to chapters 3 and 5 respectively. Finally, appendix D contains verbatim copies of Refs. [18, 84]. These two papers were produced during the author's PhD candidature, but were not related to the main theme of the thesis.

Background Theory

In this chapter we describe some basic theoretical background followed by the specific theoretical methods used in the thesis. We start with some elementary remarks about condensation in Bose gases and then show how nonrelativistic bosonic quantum field theory arises from single-particle quantum mechanics. We next derive the Gross-Pitaevskii equation from the variational point of view. The projected Gross-Pitaevskii equation is covered in some detail, including a discussion of the motivation and procedure for taking the classical limit.

2.1 Bose statistics and BEC in noninteracting systems

Bose-Einstein condensation is a unique phase transition because it exists even in the absence of interactions between particles: it is a purely statistical effect. Without interactions, the behaviour of a Bose gas can readily be analysed using the methods of elementary statistical mechanics. The following outlines the argument, showing that the occupation of the lowest energy state is extensive below some transition temperature, which we then compute.

2.1.1 The Bose distribution

In this section we derive the Bose distribution from basic statistical mechanical considerations. Before we start it is worth noting that the only quantum mechanical ingredient required is indistinguishability of particles. When we talk about microstates of the system, we may think about either classical states with a particular total energy and number of particles, or quantum eigenstates of the Hamiltonian.

In the grand canonical ensemble, the probability of observing a given microstate s of a system is proportional to the Gibbs factor:

$$P(s) \propto e^{-\beta[E(s) - \mu N(s)]}. \quad (2.1)$$

Here $E(s)$ is the energy, $N(s)$ is the number of particles in the microstate and $\beta = 1/k_B T$ is the inverse temperature. The Gibbs factor comes about by considering a system that can exchange energy and particles with a much larger reservoir, and applying the fundamental assumption of statistical mechanics — that all microstates in a closed system are equally likely — to the closed combination of system and reservoir (see, for example, [115, §7.1]). The normalisation constant for P is the sum of the Gibbs factors for all states, known as the partition function

$$\mathcal{Z} = \sum_s e^{-\beta[E(s) - \mu N(s)]}, \quad (2.2)$$

with this definition we have $P(s) = \mathcal{Z}^{-1} e^{-\beta[E(s) - \mu N(s)]}$.

We consider a system with a number of discrete single-particle quantum states or *modes*, $\{|0\rangle, |1\rangle, |2\rangle, \dots\}$. Letting the single-particle energies for the modes be ϵ_i and considering the case where particles have no interaction energy, the total energy is $E(s) = \sum_i \epsilon_i n_i(s)$, with corresponding total number $N(s) = \sum_i n_i(s)$ where n_i is the occupation for the i th mode. With these expressions, the partition function for a noninteracting system is

$$\mathcal{Z} = \sum_s e^{-\beta \sum_i (\epsilon_i - \mu) n_i(s)} \quad (2.3)$$

$$= \sum_s \prod_i e^{-\beta (\epsilon_i - \mu) n_i(s)} \quad (2.4)$$

where the product ranges over all modes.

To proceed further, we need to make use of Bose statistics to specify over which states the sum \sum_s occurs. Because bosons are indistinguishable, a state is uniquely defined by the number of particles in each mode, $s = [n_1(s), n_2(s), \dots]$. Therefore, after some thought¹ \mathcal{Z} can be rewritten in the simpler form

$$\mathcal{Z} = \prod_i \sum_{n=0}^{\infty} e^{-\beta (\epsilon_i - \mu) n} \quad (2.5)$$

$$\equiv \prod_i \mathcal{Z}_i, \quad (2.6)$$

¹To make this transformation slightly less mysterious, consider only two modes. Enumerating the possible states as $[0, 0], [1, 0], [0, 1], [2, 0], [1, 1], [0, 2], [3, 0], [2, 1], [1, 2], [0, 3], \dots$, and abbreviating $a = e^{-\beta(\epsilon_1 - \mu)}$, $b = e^{-\beta(\epsilon_2 - \mu)}$, we see that we're dealing with the expression

$$\begin{aligned} \mathcal{Z} &= \sum_s a^{n_1(s)} b^{n_2(s)} = 1 + a + b + a^2 + ab + b^2 + a^3 + a^2b + ab^2 + b^3 + \dots \\ &= (1 + a + a^2 + a^3 + \dots)(1 + b + b^2 + b^3 + \dots) = \sum_{n=0}^{\infty} a^n \sum_{n=0}^{\infty} b^n. \end{aligned}$$

where we have defined the single-mode partition function $\mathcal{Z}_i \equiv \sum_{n=0}^{\infty} e^{-\beta(\epsilon_i - \mu)n}$. We note that the simplicity of Eq. (2.5) is what makes the analytical calculation tractable, and is peculiar to the grand canonical ensemble.

An explicit formula for \mathcal{Z}_i is easily obtained by summing the geometric series:

$$\mathcal{Z}_i = \sum_{n=0}^{\infty} e^{-\beta(\epsilon_i - \mu)n} = \sum_{n=0}^{\infty} \left(e^{-\beta(\epsilon_i - \mu)} \right)^n = \frac{1}{1 - e^{-\beta(\epsilon_i - \mu)}}. \quad (2.7)$$

This puts us in a position to calculate the expected occupation of the j th mode. Because of the factorised form of Eq. (2.5), we can show that functions of n_j are normalised using only \mathcal{Z}_j rather than \mathcal{Z} ,

$$\langle f(n_j) \rangle = \sum_s f(n_j(s)) P(s) = \frac{1}{\mathcal{Z}_j} \sum_{n=0}^{\infty} f(n) e^{-\beta(\epsilon_j - \mu)n} \quad (2.8)$$

for any function f . For the particular case of $f(n) = n$ we can evaluate the sum², leading to the famous Bose distribution function

$$\langle n_j \rangle = \frac{1}{e^{\beta(\epsilon_j - \mu)} - 1}. \quad (2.9)$$

2.1.2 Condensation in the 3D homogeneous Bose gas

Having derived the Bose distribution, we turn our attention to the statistical mechanics of Bose condensation in a 3D gas. The expected total number of atoms in the system is given by the sum over modes

$$N = \sum_i \langle n_i \rangle = \sum_i \frac{1}{e^{\beta(\epsilon_i - \mu)} - 1}. \quad (2.10)$$

For a given system the modal energies $\{\epsilon_i\}$ are known, so at a particular total number N and inverse temperature β we may in principle use this equation to determine μ . However, solving this equation is not possible except in the simplest cases, and we therefore turn to approximations.

The general procedure is to split the atoms into two groups: the expected numbers in the ground and excited states, N_0 and N_{ex} respectively. The original argument due to Einstein is that in the 3D homogeneous gas there is an upper bound on the number of excited state atoms N_{ex} at any given temperature. For any system with more atoms than this bound, the

²These kinds of sums may be reduced to the sum of a geometric series using the common derivative trick:

$$\sum_{n=0}^{\infty} n x^n = \sum_{n=0}^{\infty} n e^{\ln(x)n} = \frac{\partial}{\partial y} \sum_{n=0}^{\infty} e^{y n} \Big|_{y=\ln x} = \frac{\partial}{\partial y} \frac{1}{1 - e^y} \Big|_{y=\ln x} = \frac{x}{(1 - x)^2}.$$

remaining atoms must be found condensed in the ground state [37]:

[...] something similar happens as when isothermally compressing a vapour beyond the volume of saturation. A separation occurs; a part “condenses”, the rest remains a “saturated ideal gas”.

The expected number of excited state atoms is the sum over excited state modal occupations,

$$N_{\text{ex}} = \sum_{i \neq 0} \frac{1}{e^{\beta(\epsilon_i - \mu)} - 1}. \quad (2.11)$$

To estimate N_{ex} analytically we need a few transformations and approximations, starting by rewriting the sum as a sum over the mode energies rather than state indices. Defining the discrete density of states $g_{\text{dsc}}(\epsilon)$ to be the number of modes at energy ϵ , we have

$$N_{\text{ex}} = \sum_{\epsilon} g_{\text{dsc}}(\epsilon) \frac{1}{e^{\beta(\epsilon - \mu)} - 1}. \quad (2.12)$$

Next, we recognise that $N_0 > 0$ implies $\mu < \epsilon_0$, which in turn implies $1/(e^{\beta(\epsilon_i - \mu)} - 1) < 1/(e^{\beta(\epsilon_i - \epsilon_0)} - 1)$ so that we can replace μ with ϵ_0 to bound N_{ex} from above:

$$N_{\text{ex}} < \sum_{\epsilon} g_{\text{dsc}}(\epsilon) \frac{1}{e^{\beta(\epsilon - \epsilon_0)} - 1}. \quad (2.13)$$

In the particular case of the ideal gas, $\epsilon_0 = 0$ and we have

$$N_{\text{ex}} < \sum_{\epsilon} g_{\text{dsc}}(\epsilon) \frac{1}{e^{\beta\epsilon} - 1}. \quad (2.14)$$

Unfortunately evaluating this sum is not straightforward, so to progress further we transform it into an integral via an approximation. We have

$$N_{\text{ex}} \lesssim \int d\epsilon g(\epsilon) \frac{1}{e^{\beta(\epsilon - \mu)} - 1}, \quad (2.15)$$

where $g(\epsilon)$ is the density of states that will correspond to a smoothed version of $g_{\text{dsc}}(\epsilon)$.

The most elementary method for computing an appropriately smoothed g is to first evaluate the total number of states $G(\epsilon)$ with energy less than ϵ (see, for example, [96]). With the simplest approximate evaluation, $G(\epsilon)$ turns out to be smooth and we may define a smooth g using $g(\epsilon) \equiv \frac{dG}{d\epsilon}$. For the 3D gas in a finite periodic box of size $L \times L \times L$, the possible modes

are indexed by the integer wavenumber $\mathbf{q} \in \mathbb{Z}^3$ and have energy

$$E(\mathbf{q}) = \frac{2\hbar^2\pi^2\mathbf{q}^2}{mL^2}. \quad (2.16)$$

The number of modes with energy less than ϵ is then approximately the volume of the sphere of radius $\sqrt{mL^2\epsilon/2\hbar^2\pi^2}$ containing the set $\{\mathbf{q} \in \mathbb{R}^3: E(\mathbf{q}) < \epsilon\}$, that is,

$$G(\epsilon) = \frac{4\pi}{3} \left(\frac{mL^2\epsilon}{2\hbar^2\pi^2} \right)^{3/2}. \quad (2.17)$$

Differentiating G leads to the density of states

$$g(\epsilon) = \frac{dG}{d\epsilon} = \frac{4\sqrt{2}\pi L^3 m^{3/2}}{h^3} \epsilon^{1/2}, \quad (2.18)$$

which we see varies as the square root of ϵ .

Putting this into Eq. (2.15) and performing the integral³, we finally arrive at

$$N_{\text{ex}} \lesssim \left(\frac{2\pi mL^2 k_B}{h^2} \right)^{3/2} \zeta\left(\frac{3}{2}\right) T^{3/2} \quad (2.19)$$

where $\zeta(n) = \sum_{k=1}^{\infty} 1/k^n$ is the Riemann zeta function, with $\zeta(\frac{3}{2}) \approx 2.61$. This expression tells us that there must be some critical temperature T_c below which N is strictly greater than N_{ex} , and the remaining $N - N_{\text{ex}}$ atoms must go into the ground state. At the critical temperature, $N = N_{\text{ex}}(T_c)$ which implies

$$T_c = \frac{h^2 N^{2/3}}{2\pi mL^2 k_B \zeta\left(\frac{3}{2}\right)^{2/3}}. \quad (2.20)$$

The fraction of atoms in the ground state as a function of temperature is

$$f_c = \frac{N - N_{\text{ex}}}{N} = 1 - \left(\frac{T}{T_c} \right)^{3/2}. \quad (2.21)$$

Determining the condensate fraction using the standard derivation presented above depends critically on the functional form of the density of states. In turn, the density of states depends on the dimensionality and any trapping potential, both of which modify the possible single-particle

³The integral may be evaluated by expanding in powers of e^{-x} , followed by a change of variables and recognising the Riemann zeta function ζ and gamma function Γ :

$$\int_0^{\infty} dx \frac{x^{1/2}}{e^x - 1} = \sum_{n=1}^{\infty} \int_0^{\infty} dx x^{1/2} e^{-xn} = \sum_{n=1}^{\infty} n^{-3/2} \int_0^{\infty} dx x^{1/2} e^{-x} = \zeta\left(\frac{3}{2}\right) \Gamma\left(\frac{3}{2}\right).$$

quantum states. Indeed, in the two-dimensional homogeneous gas the integral in Eq. (2.15) fails to converge, suggesting — though not by itself proving — an absence of condensation.

2.2 Field theory and interacting Bose gases

In this section we develop some basics of quantum field theory in a form that applies to cold gases of atomic bosons. In particular, the theory developed below is strictly nonrelativistic, and assumes bosonic symmetry for the many-particle quantum state.

2.2.1 Second quantisation

The formalism of *second quantisation*⁴ is an elegant and compact way to express the quantum mechanics of identical particles. Second quantisation succeeds by incorporating the symmetries of the many-particle state into the theory from the start. At heart, it may be viewed as a convenient reformulation of basic many-particle quantum mechanics, extended to allow for superpositions of states with different particle number. There are many good explanations of second quantisation in the literature; for additional exposition we refer the reader to Refs. [2, 122].

Many-particle symmetries and the number state basis

To describe second quantisation we start with the quantum mechanics of an N particle system. In the general case, the Hilbert space for a system of N distinguishable particles is the tensor product of N one-particle Hilbert spaces (see, for example, [93, §2.2.8]):

$$\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \cdots \otimes \mathcal{H}_N. \quad (2.22)$$

Accordingly, the possible states of N *identical* particles live in the Hilbert space

$$\mathcal{H} = \underbrace{\mathcal{H}_1 \otimes \mathcal{H}_1 \otimes \cdots \otimes \mathcal{H}_1}_{N \text{ times}} \equiv \mathcal{H}_1^N. \quad (2.23)$$

States from the natural tensor product basis for \mathcal{H}_1^N have the form⁵ $|j_1\rangle|j_2\rangle\cdots|j_N\rangle$, where the j_i are positive integers and we use $\{|1\rangle, |2\rangle, \dots\}$ as a basis for \mathcal{H}_1 .

⁴The name “second quantisation” is somewhat obscure. Historically, one method of deriving a non-interacting quantum field theory was to apply the procedure of canonical quantisation to the Schrödinger field (see [122, page 81] for a discussion). Insofar as one step of quantisation has already been used to obtain the Schrödinger equation, it appeared that this procedure constituted quantising a second time.

⁵We omit the tensor product symbols between kets for brevity; $|a\rangle|b\rangle \equiv |a\rangle \otimes |b\rangle$.

It is possible to solve problems directly in \mathcal{H}_1^N . However, the indistinguishability of particles implies a strict symmetry requirement on which states of \mathcal{H}_1^N correspond to realisable physical states. In the case of bosons, the state must remain unchanged under an interchange of particle indices. For example, a two-particle state $|\psi\rangle = \sum_{i,j} c_{ij}|i\rangle|j\rangle$ must have the property $c_{ij} = c_{ji}$; in a continuous basis this takes the more familiar form $\psi(x_1, x_2) = \psi(x_2, x_1)$. We note that the correct derivation of the allowable boson and fermion symmetry rules is considerably more subtle than the common argument based on swapping labels in the many-particle wavefunction [80].

Taking symmetry into account, the general state for N bosons may be written

$$|j_1, \dots, j_N\rangle_{\text{sym}} = \frac{1}{\sqrt{N! \prod_{i=1}^{\infty} n_i!}} \sum_P |j_{P_1}\rangle |j_{P_2}\rangle \cdots |j_{P_N}\rangle \quad (2.24)$$

where the sum is over all permutations P of the integers $\{1, \dots, N\}$ and n_i is the number of particles in mode $|i\rangle$ as before. Such wavefunctions are usually extremely inconvenient to use for computational purposes due to the large number $\prod_{i=1}^{\infty} n_i! / N!$ of unique terms.

At this stage it is useful to introduce the number state notation. To ensure uniqueness, the indices in the state $|j_1, \dots, j_N\rangle_{\text{sym}}$ must be ordered — for example, $|1, 1, 1, 2\rangle_{\text{sym}}$ and $|1, 1, 2, 1\rangle_{\text{sym}}$ are the same state due to the sum over all permutations. This redundancy can be removed by instead listing the occupations of all modes. We have

$$|n_1, n_2, \dots\rangle \equiv |\underbrace{1, \dots, 1}_{n_1 \text{ times}}, \underbrace{2, \dots, 2}_{n_2 \text{ times}}, \dots\rangle_{\text{sym}}, \quad (2.25)$$

so that, for example, $|1, 1, 1, 2\rangle_{\text{sym}} = |3, 1, 0, 0, \dots\rangle$. The state $|n_1, n_2, \dots\rangle$ has a well-defined number of bosons in each mode and is therefore called a *number state*.

Fock space and the creation and annihilation operators

Let \mathcal{F}^N be the subspace of \mathcal{H}_1^N spanned by all the correctly symmetrised number states. \mathcal{F}^N contains all the possible states of the N bosons, and therefore avoids the redundancy of non-physical states contained in \mathcal{H}_1^N . The *Fock space* is constructed via the direct sum⁶ of all the \mathcal{F}^N with varying N :

$$\mathcal{F} = \bigoplus_{N=0}^{\infty} \mathcal{F}^N. \quad (2.26)$$

⁶The direct sum of two Hilbert spaces, $\mathcal{H}_1 \oplus \mathcal{H}_2$, is the space of ordered pairs $|(\psi, \phi)\rangle \in \mathcal{H}_1 \times \mathcal{H}_2$, with inner product defined by

$$\langle(\psi_1, \phi_1)|(\psi_2, \phi_2)\rangle = \langle\psi_1|\psi_2\rangle + \langle\phi_1|\phi_2\rangle.$$

Note the presence of the Hilbert space \mathcal{F}^0 of zero particles which has a single basis vector known as the *vacuum state* and conventionally written $|0\rangle$.

The efficiency of second quantisation arises because we can perform calculations entirely within \mathcal{F} and the symmetry of the number states is built in from the start. The formalism of *creation* and *annihilation operators* is the tool used for this task. We define the creation operator for the i th mode by its action on the number basis:

$$\hat{a}_i^\dagger |n_1, \dots, n_i, \dots\rangle = (n_i + 1)^{1/2} |n_1, \dots, n_i + 1, \dots\rangle. \quad (2.27)$$

That is, the action of the creation operator \hat{a}_i^\dagger is to create a particle in the i th mode. Note that this is a complete specification of \hat{a}_i^\dagger , because we have defined its action for every number state and the number states are a basis for \mathcal{F} . From the definition we have $\langle m + 1 | \hat{a}_i^\dagger | n \rangle = (n + 1)^{1/2} \delta_{mn}$; taking the conjugate and rearranging shows that the adjoint \hat{a}_i is the annihilation operator for mode i because it reduces the occupation by one:

$$\hat{a}_i |n_1, \dots, n_i, \dots\rangle = n_i^{1/2} |n_1, \dots, n_i - 1, \dots\rangle. \quad (2.28)$$

A basic and important property of the creation and annihilation operators is their commutation relations. These can be computed directly from Eqs. (2.27) and (2.28) yielding

$$[\hat{a}_i, \hat{a}_j^\dagger] = \delta_{ij}, \quad [\hat{a}_i^\dagger, \hat{a}_j^\dagger] = 0, \quad [\hat{a}_i, \hat{a}_j] = 0. \quad (2.29)$$

The commutation relations are vital for calculations, but may also be used as an alternate starting point to Eq. (2.27), resulting in the same operators.

The operators \hat{a}_i^\dagger are defined with respect to a particular single-particle basis $\{|i\rangle\}$; a different set of operators \hat{b}_κ^\dagger arises if we consider a different basis $\{|\kappa\rangle\}$. (Note that we are abusing the notation somewhat and treating kets with Greek indices as a distinct basis from those with Latin indices.) The change of basis law may be derived by considering the action of \hat{a}_i^\dagger on the vacuum state, and using the resolution of identity $\hat{1} = \sum_\kappa |\kappa\rangle \langle \kappa|$ in the one-particle Hilbert space:

$$\hat{a}_i^\dagger |0\rangle = |i\rangle = \left(\sum_\kappa |\kappa\rangle \langle \kappa| \right) |i\rangle = \sum_\kappa \langle \kappa | i \rangle |\kappa\rangle = \sum_\kappa \langle \kappa | i \rangle \hat{b}_\kappa^\dagger |0\rangle. \quad (2.30)$$

The change of basis is therefore given by

$$\hat{a}_i^\dagger = \sum_\kappa \langle \kappa | i \rangle \hat{b}_\kappa^\dagger \quad \text{and} \quad \hat{a}_i = \sum_\kappa \langle i | \kappa \rangle \hat{b}_\kappa. \quad (2.31)$$

Using the analogous transformation law for the continuous position basis $\{|\mathbf{x}\rangle\}$ gives us the powerful concept of the *boson field operator*

$$\hat{\psi}(\mathbf{x}) = \sum_i \langle \mathbf{x} | i \rangle \hat{a}_i = \sum_i \psi_i(\mathbf{x}) \hat{a}_i \quad (2.32)$$

where $\psi_i(\mathbf{x}) \equiv \langle \mathbf{x} | i \rangle$ is the shape of mode i in position space. Note that we use a Greek letter for the field operator by convention; we could equally well have used $\hat{a}^\dagger(\mathbf{x})$ to emphasise the similarity with the operators \hat{a}_i^\dagger for the discrete basis. We interpret the operation of the object $\hat{\psi}^\dagger(\mathbf{x})$ on a quantum state as the creation of a particle at position \mathbf{x} . Alternatively, it may be thought of as a field that associates an operator on \mathcal{F} to each point of space. The commutation relations for the continuous case are analogous to the discrete case:

$$[\hat{\psi}(\mathbf{x}), \hat{\psi}^\dagger(\mathbf{y})] = \delta(\mathbf{x} - \mathbf{y}), \quad [\hat{\psi}^\dagger(\mathbf{x}), \hat{\psi}^\dagger(\mathbf{y})] = 0, \quad [\hat{\psi}(\mathbf{x}), \hat{\psi}(\mathbf{y})] = 0. \quad (2.33)$$

Various expressions are shorter and more familiar when written in the position basis using $\hat{\psi}$, as will become clear in the following sections.

Representation of operators on Fock space

The creation and annihilation operators are useful because of two important properties: First, they act in a particularly straightforward way on the number state basis of \mathcal{F} , as per the definition. Second, the operators of interest on \mathcal{H}_1^N can be written simply in terms of \hat{a}_i^\dagger and \hat{a}_i . Once the necessary operators are expressed in the natural basis for \mathcal{F} there is no longer any need to consider explicitly symmetrised states defined on \mathcal{H}_1^N . This section examines the representation of one and two-particle operators on Fock space, largely following the development in Ref. [2].

We examine the case of single-particle operators first. Consider an operator \hat{q} acting on \mathcal{H}_1 . \hat{q} may be extended to act on the space of N particles as the sum

$$\hat{Q} = \sum_{i=1}^N \hat{Q}_i \quad (2.34)$$

where \hat{Q}_i represents \hat{q} operating on the i th particle⁷. Suppose now that \hat{q} is diagonal in some basis $\{|\kappa\rangle\}$ of \mathcal{H}_1 so that $\hat{Q} = \sum_\kappa q_{\kappa\kappa} |\kappa\rangle \langle \kappa|$. Using the explicitly symmetrised wavefunction in

⁷ \hat{Q}_i may be written explicitly as $\hat{Q}_i = \hat{\mathbb{1}}^{\otimes(N-i)} \otimes \hat{q} \otimes \hat{\mathbb{1}}^{\otimes i}$ where $\hat{\mathbb{1}}^{\otimes i} \equiv \underbrace{\hat{\mathbb{1}} \otimes \cdots \otimes \hat{\mathbb{1}}}_{i \text{ times}}$ is the identity on \mathcal{H}_1^i , constructed from i copies of the identity operator $\hat{\mathbb{1}}$ on \mathcal{H}_1 .

Eq. (2.24) the action of \hat{Q} on a number state may be computed. After some algebra, we find

$$\hat{Q}|n_{\kappa 1}, n_{\kappa 2}, \dots\rangle = \left(\sum_{\kappa} q_{\kappa\kappa} n_{\kappa}\right)|n_{\kappa 1}, n_{\kappa 2}, \dots\rangle \quad (2.35)$$

$$= \left(\sum_{\kappa} q_{\kappa\kappa} \hat{b}_{\kappa}^{\dagger} \hat{b}_{\kappa}\right)|n_{\kappa 1}, n_{\kappa 2}, \dots\rangle \quad (2.36)$$

$$\implies \hat{Q} = \sum_{\kappa} q_{\kappa\kappa} \hat{b}_{\kappa}^{\dagger} \hat{b}_{\kappa} = \sum_{\kappa} \langle \kappa | \hat{q} | \kappa \rangle \hat{b}_{\kappa}^{\dagger} \hat{b}_{\kappa} \quad (2.37)$$

In the second line we have used the fact that the *occupation number operator* for the κ th mode is $\hat{n}_{\kappa} = \hat{b}_{\kappa}^{\dagger} \hat{b}_{\kappa}$. Changing to the non-diagonal basis using Eq. (2.31) gives the general expression for any one-particle operator

$$\hat{Q} = \sum_{i,j} \langle i | \hat{q} | j \rangle \hat{a}_i^{\dagger} \hat{a}_j. \quad (2.38)$$

A related procedure can be carried out for multi-particle operators, with similar results. If \hat{q}' is a two-particle operator acting on the space \mathcal{H}_1^2 , an extension to \mathcal{H}_1^N is

$$\hat{Q}' = \sum_{\substack{i,j=1 \\ i \neq j}}^N \hat{Q}'_{ij} \quad (2.39)$$

where \hat{Q}'_{ij} is \hat{q}' acting on the subspace of the i th and j th particles. The operator \hat{Q}' can be expressed in terms of the creation and annihilation operators. The procedure is similar to the derivation for the single-particle operators above, with some care required in handling the sum so that $i \neq j$. The general result is

$$\hat{Q}' = \sum_{i,j,k,l} (\langle i | \otimes \langle j |) \hat{q}' (|k\rangle \otimes |l\rangle) \hat{a}_i^{\dagger} \hat{a}_j^{\dagger} \hat{a}_k \hat{a}_l. \quad (2.40)$$

Note that including both orderings (i, j) and (j, i) in Eq. (2.39) is redundant because \hat{q}' is symmetric for identical particles so $\hat{Q}'_{ij} = \hat{Q}'_{ji}$.

The second-quantised operators we will be interested in are conveniently compact and familiar in the position basis, as exemplified by the Hamiltonian. The single-particle part of the Hamiltonian⁸,

$$H_{\text{sp}} = -\frac{\hbar^2}{2m} \nabla_{\mathbf{x}}^2 + V(\mathbf{x}), \quad (2.41)$$

⁸It is conventional to omit the hat from H_{sp} and other operators acting on the single-particle configuration space. This is somewhat inconsistent, but serves as a useful reminder of the subspace on which these operators act when dealing with the full Fock space.

extended to Fock space is

$$\hat{H}_1 = \int d\mathbf{x} \, \hat{\psi}^\dagger(\mathbf{x}) \left[-\frac{\hbar^2}{2m} \nabla_{\mathbf{x}}^2 + V(\mathbf{x}) \right] \hat{\psi}(\mathbf{x}). \quad (2.42)$$

It is notable that this equation has only a single integral over the variable \mathbf{x} rather than the double integral — or sum in the discrete case — which would be expected in general (cf. Eq. (2.38)). This is because H_{sp} acts only locally in the position basis, although it is not quite diagonal.

If the interaction potential between a pair of particles located at \mathbf{x} and \mathbf{x}' is given by $V_{\text{int}}(\mathbf{x} - \mathbf{x}')$, the interaction term in the N particle Hamiltonian is

$$\sum_{\substack{i,j=1 \\ i < j}}^N V_{\text{int}}(\mathbf{x}_i - \mathbf{x}_j) = \frac{1}{2} \sum_{\substack{i,j=1 \\ i \neq j}}^N V_{\text{int}}(\mathbf{x}_i - \mathbf{x}_j). \quad (2.43)$$

This is diagonal in position space, so writing it in second-quantised form is simple:

$$\hat{H}_2 = \frac{1}{2} \iint d\mathbf{x} \, d\mathbf{x}' \, \hat{\psi}^\dagger(\mathbf{x}) \hat{\psi}^\dagger(\mathbf{x}') V_{\text{int}}(\mathbf{x} - \mathbf{x}') \hat{\psi}(\mathbf{x}') \hat{\psi}(\mathbf{x}). \quad (2.44)$$

To summarise, the second-quantised Hamiltonian for the dilute Bose gas is given by

$$\hat{H} = \int d\mathbf{x} \, \hat{\psi}^\dagger(\mathbf{x}) H_{\text{sp}} \hat{\psi}(\mathbf{x}) + \frac{1}{2} \iint d\mathbf{x} \, d\mathbf{x}' \, \hat{\psi}^\dagger(\mathbf{x}) \hat{\psi}^\dagger(\mathbf{x}') V_{\text{int}}(\mathbf{x} - \mathbf{x}') \hat{\psi}(\mathbf{x}') \hat{\psi}(\mathbf{x}). \quad (2.45)$$

This Hamiltonian is the starting point for every theoretical analysis of the single species Bose gas [107]; it contains a complete description of the physics in the low-energy regime.

2.2.2 The low-energy Hamiltonian

Experiments with ultracold gases often take place in the dilute regime — the typical interatomic distance is much larger than the range of the potential V_{int} . As a consequence, the interactions may be treated as two-body scattering events to a high degree of accuracy. Solving the two-body scattering problem involves expanding the wavefunction in a series of partial waves. At sufficiently low energies only the s partial wave is important and the scattering may be characterised by a single parameter a_s called the *s-wave scattering length*. The scattering length is on the order of $100a_0$ for alkali atoms, where a_0 is the Bohr radius, though there is considerable experimental flexibility in tuning the value using Feshbach resonances. For a detailed account of the scattering theory and related issues, we direct the reader to Ref. [96, Ch. 5].

In practice, the dilute and low energy nature of the system implies that the theory can be greatly simplified by formally replacing the full interaction potential by a delta function

$$V_{\text{int}}(\mathbf{x} - \mathbf{x}') \rightarrow U_0 \delta(\mathbf{x} - \mathbf{x}') \quad (2.46)$$

in the Hamiltonian, where $U_0 = 4\pi\hbar^2 a_s/m$. Strictly speaking, this replacement arises by integrating out the high energy states — resulting in a two-body T-matrix description — followed by taking the zero energy limit (see, for example, [14, §2.1]). The resulting theory necessarily contains a high energy cutoff that prevents the delta potential from unphysically scattering waves of arbitrarily high momentum. Nevertheless, we ignore this complication in what follows because the main approximations used later in the thesis are not sensitive to it.

Performing the replacement of Eq. (2.46) leads to the approximate low-energy Hamiltonian:

$$\hat{H} = \int d\mathbf{x} \hat{\psi}^\dagger(\mathbf{x}) H_{\text{sp}} \hat{\psi}(\mathbf{x}) + \frac{U_0}{2} \int d\mathbf{x} \hat{\psi}^\dagger(\mathbf{x}) \hat{\psi}^\dagger(\mathbf{x}) \hat{\psi}(\mathbf{x}) \hat{\psi}(\mathbf{x}). \quad (2.47)$$

A theory built from this Hamiltonian retains all the smooth long wavelength behaviour, while all short range correlations between particles are discarded.

2.2.3 Operator equations of motion

In section 2.2.1 we wrote operators and states for N particle quantum mechanics in the convenient notation of second quantisation but we said nothing about time evolution. Clearly we may use the *Schrödinger picture* in which the quantum state on Fock space evolves according to the usual Schrödinger equation. Equivalently we may use the *Heisenberg picture* of time evolution where the operators evolve and the quantum state is fixed. The Heisenberg equation of motion for an operator \hat{A}_H is

$$i\hbar \frac{\partial \hat{A}_H(t)}{\partial t} = [\hat{A}_H(t), \hat{H}], \quad (2.48)$$

where we have assumed that the corresponding Schrödinger picture operator $\hat{A}_S = \hat{A}_H(0)$ is time independent. Note that this equation is linear — in the usual sense that linear combinations of solutions are new solutions — because the so-called superoperator $[\cdot, \hat{H}]$ is linear.

We can in principle write any desired Heisenberg picture observable in terms of the Heisenberg picture field operator $\hat{\psi}_H(\mathbf{x}, t)$. Therefore the full dynamics of the system are encoded in the behaviour of $\hat{\psi}_H(\mathbf{x}, t)$, the general evolution of which is given by

$$i\hbar \frac{\partial \hat{\psi}_H(\mathbf{x}, t)}{\partial t} = [\hat{\psi}_H(\mathbf{x}, t), \hat{H}]. \quad (2.49)$$

We now drop the subscripts for brevity, assuming that from now on we will work in the Heisenberg picture unless otherwise specified. Using the Hamiltonian in Eq. (2.47) along with the commutation relations for $\hat{\psi}^\dagger$ and $\hat{\psi}$, we obtain the operator equation of motion for $\hat{\psi}$ (see, for example, [107]):

$$i\hbar \frac{\partial \hat{\psi}(\mathbf{x}, t)}{\partial t} = H_{\text{sp}} \hat{\psi}(\mathbf{x}, t) + U_0 \hat{\psi}^\dagger(\mathbf{x}, t) \hat{\psi}(\mathbf{x}, t) \hat{\psi}(\mathbf{x}, t). \quad (2.50)$$

The apparently nonlinear form of this equation presents a puzzle when compared to the manifestly linear form of the equation of motion (2.48) for a general operator. What we have here is a tradeoff: We have removed the explicit dependence on the Hamiltonian operator so that the equation is expressed purely in terms of $\hat{\psi}$; the price is that to actually evaluate the time evolution requires computing a nonlinear function of field operators. Nevertheless, the underlying linearity of the solutions is preserved⁹.

The Heisenberg equations of motion arising from very simple Hamiltonians can sometimes be solved directly. However, including interactions as in Eq. (2.50) generally makes the problem analytically intractable and attacking it directly using numerics is also out of the question, due to the infinite number of degrees of freedom. Nevertheless, the Heisenberg equations of motion for the field operator give us a useful starting point for certain types of approximations, including the projected Gross-Pitaevskii equation discussed in section 2.4.

2.3 The Gross-Pitaevskii equation

The Gross-Pitaevskii equation¹⁰ (GPE) is a remarkably successful simplification of the full quantum field equations at the level of mean-field theory. We have already seen that for the ideal Bose gas all particles condense into a single quantum mode at zero temperature. The central approximation of the Gross-Pitaevskii theory is to assume this wavefunction is also reasonable for *dynamics* with small but nonzero interactions and excitation energies. That is,

⁹For a very simple example, consider a single-mode system with Hamiltonian $\hat{H} = \frac{1}{2} \hat{a}^\dagger \hat{a}^\dagger \hat{a} \hat{a}$ representing a two-particle interaction. The Heisenberg equation of motion for \hat{a} is

$$i\hbar \frac{\partial \hat{a}}{\partial t} = \hat{a}^\dagger \hat{a} \hat{a},$$

which appears to be nonlinear. However, \hat{H} may also be expressed in the number basis as $\hat{H} = \frac{1}{2} \sum_n n(n-1) |n\rangle \langle n|$. Writing $\hat{a}(t) = \sum_{n,m} a_{nm}(t) |n\rangle \langle m|$, the evolution equation for the components of $\hat{a}(t)$ is clearly linear:

$$i\hbar \frac{\partial a_{nm}}{\partial t} = \frac{1}{2} [m(m-1) - n(n-1)] a_{nm}.$$

¹⁰The GPE is commonly known as the (cubic) nonlinear Schrödinger equation (NLSE) in other areas of physics and in mathematics.

we assume the Schrödinger picture wavefunction is

$$\Phi(\mathbf{x}_1, \dots, \mathbf{x}_N, t) \equiv \prod_{i=1}^N \varphi(\mathbf{x}_i, t), \quad (2.51)$$

and compute an equation of motion for φ so that the dynamics of Φ approximates the full quantum dynamics of Ψ as closely as possible. The assumed form for the wavefunction given above is the *time dependent Hartree-Fock* ansatz. (We note that N bosons in a single mode is a special case where the fully symmetrised Hartree-Fock ansatz is identical to the simple product state, also known as the Hartree ansatz.)

There are numerous methods for deriving the equations of motion for φ . In the cold atoms literature the most commonly used are geared toward computing not only an equation for φ but also the next order corrections, giving the so-called Hartree-Fock-Bogoliubov theories. These techniques may be broadly grouped into the *symmetry breaking* (see, for example, [24]) and *number conserving* formalisms (see, for example, [48, 88, 49]). For a discussion of the difference between these techniques, see Refs. [88, 107].

To derive only the lowest order theory, we choose a more elementary method based on a time dependent variational principle.

2.3.1 The Dirac-Frenkel time dependent variational principle

The Dirac-Frenkel time dependent variational principle [81, §II.1] has been commonly used in quantum chemistry and nuclear theory in the context of time dependent Hartree-Fock theory [70]. It is also an expedient method for deriving the time dependent GPE [96, § 7.1] and with that in mind we briefly outline some of its generic properties before making use of it in the next section. The Dirac-Frenkel variational principle is characterised by stationarity of the quantum effective action¹¹

$$S[\Psi] = \int_{t_0}^{t_1} dt \langle \Psi | i\hbar \partial_t - \hat{H} | \Psi \rangle = \int_{t_0}^{t_1} dt \int d\mathbf{x} \Psi^* (i\hbar \partial_t - \hat{H}) \Psi \quad (2.52)$$

with respect to small variations in Ψ^* . For an unconstrained state Ψ in the full Hilbert space \mathcal{H} , the variational principle is equivalent to

$$\frac{\delta S}{\delta \Psi^*} = 0, \quad (2.53)$$

¹¹In path integral quantisation, one builds the quantum theory out of a classical action integrated over all paths. This is distinct from the full quantum “effective action” for the same system, which is a functional of the quantum rather than classical state.

where $\delta S/\delta\Psi^*$ is the functional derivative of the nonholomorphic function¹² S with respect to the complex variable Ψ^* . For a detailed discussion including how to compute with these derivatives in practice, see appendix A.2. Evaluating the functional derivative gives

$$\frac{\delta S}{\delta\Psi^*} = (i\hbar\partial_t - \hat{H})\Psi, \quad (2.54)$$

which immediately yields the time dependent Schrödinger equation when set to zero.

In contrast to the trivial case above, the variational principle is most useful when constraining Ψ to some submanifold \mathcal{M} of \mathcal{H} ; in our case \mathcal{M} will be the set of all states of the form given in Eq. (2.51). In fact, the Dirac-Frenkel variational principle arises as a result of trying to find the “best possible” approximation of this kind: Consider some *approximate* evolution $\Psi(t) \in \mathcal{M}$. According to the Schrödinger equation, the full dynamics attempts to evolve $\Psi(t)$ in the direction $\hat{H}\Psi(t)$, but this evolution generally takes us out of the manifold \mathcal{M} and into the bulk of the higher dimensional space \mathcal{H} . To avoid this we want an evolution direction v constrained to the tangent space $T_{\Psi(t)}\mathcal{M}$ of \mathcal{M} at $\Psi(t)$ but chosen so that the error $\|v - \hat{H}\Psi(t)\|$ is as small as possible. Evolving Ψ according to v means $i\hbar\partial_t\Psi(t) = v$ by definition, so we want to minimise $\|i\hbar\partial_t\Psi(t) - \hat{H}\Psi(t)\|$. This leads to the condition that the residual of the Schrödinger equation is perpendicular to the tangent space $T_{\Psi(t)}\mathcal{M}$:

$$\langle u | i\hbar\partial_t - \hat{H} | \Psi \rangle = 0 \quad \forall u \in T_{\Psi}\mathcal{M}. \quad (2.55)$$

The action in the Dirac-Frenkel variational principle follows from the orthogonality condition above as discussed in Ref. [81]. It is also shown that the system of equations arising from the variational principle automatically conserves the expected energy, $E[\Psi] = \langle \Psi | \hat{H} | \Psi \rangle$. Also conserved is any operator A that commutes with the Hamiltonian and keeps $A\Psi$ within $T_{\Psi}\mathcal{M}$:

$$A\Psi \in T_{\Psi}\mathcal{M} \quad \forall \Psi \in \mathcal{M} \cap D(A) \quad (2.56)$$

where $D(A)$ is the domain of A .

2.3.2 Deriving the GPE

Deriving the GPE is a straightforward application of the variational principle introduced in the previous section, with approximation manifold \mathcal{M} equal to the set of all states of the form

¹²A complex-valued function of complex arguments is *holomorphic* if it is complex differentiable according to the usual limit-based definition of derivative in the complex plane. For nonholomorphic functions we must use a modified definition of the derivative; see section A.1.

given in Eq. (2.51). For convenience, we repeat the action

$$S[\Psi] = \int_{t_0}^{t_1} dt \langle \Psi | i\hbar \partial_t - \hat{H} | \Psi \rangle. \quad (2.57)$$

To compute this for our particular assumed wavefunction Φ , we need expressions for $\langle \Phi | \hat{H} | \Phi \rangle$ and $\langle \Phi | i\hbar \partial_t | \Phi \rangle$. For the first part, we note that our wavefunction at some time t has the form $\Phi = |N, 0, \dots\rangle$ in second-quantised notation, where the only occupied mode is $\varphi(t)$. The Schrödinger picture field operators act on the state as

$$\hat{\psi} |N, 0, \dots\rangle = \varphi(\mathbf{x}, t) \sqrt{N} |N-1, 0, \dots\rangle, \quad (2.58)$$

$$\hat{\psi}^\dagger |N-1, 0, \dots\rangle = \varphi^*(\mathbf{x}, t) \sqrt{N} |N, 0, \dots\rangle \quad (2.59)$$

and therefore the expectation value of the Hamiltonian in Eq. (2.47) is

$$\begin{aligned} \langle N, 0, \dots | \hat{H} | N, 0, \dots \rangle &= N \int d\mathbf{x} \varphi^*(\mathbf{x}, t) H_{\text{sp}} \varphi(\mathbf{x}, t) \\ &+ \frac{U_0}{2} N(N-1) \int d\mathbf{x} \varphi^*(\mathbf{x}, t) \varphi^*(\mathbf{x}, t) \varphi(\mathbf{x}, t) \varphi(\mathbf{x}, t). \end{aligned} \quad (2.60)$$

Alternatively, this expression could be obtained by working directly with the sums over all particles in the first-quantised Hamiltonian¹³.

For the time derivative, it is simplest to work in the first-quantised position representation; we see that

$$i\hbar \partial_t \Phi = i\hbar \partial_t \prod_{i=1}^N \varphi(\mathbf{x}_i, t) = i\hbar \sum_{i=1}^N \partial_t \varphi(\mathbf{x}_i, t) \prod_{\substack{j=1 \\ j \neq i}}^N \varphi(\mathbf{x}_j, t), \quad (2.61)$$

so that

$$\langle \Phi | i\hbar \partial_t | \Phi \rangle = i\hbar \int d\mathbf{x}_1 \cdots \int d\mathbf{x}_N \prod_{i=1}^N \varphi^*(\mathbf{x}_i, t) \partial_t \prod_{k=1}^N \varphi(\mathbf{x}_k, t) \quad (2.62)$$

$$= Ni\hbar \int d\mathbf{x} \varphi^*(\mathbf{x}, t) \partial_t \varphi(\mathbf{x}, t). \quad (2.63)$$

¹³The first-quantised Hamiltonian is

$$H(\mathbf{x}_1, \dots, \mathbf{x}_N) = \sum_{i=1}^N H_{\text{sp}}(\mathbf{x}_i) + \frac{U_0}{2} \sum_{\substack{i,j=1 \\ i < j}}^N \delta(\mathbf{x}_i - \mathbf{x}_j),$$

which is equivalent to Eq. (2.47) when the number of particles is fixed at N .

Putting the parts of the action together, we have

$$S[\varphi] = N \int_{t_0}^{t_1} dt \int d\mathbf{x} \varphi^*(\mathbf{x}, t) [i\hbar\partial_t - H_{\text{sp}}] \varphi(\mathbf{x}, t) - \frac{U_0}{2} (N-1) |\varphi(\mathbf{x}, t)|^4. \quad (2.64)$$

To simplify further, we introduce the *condensate wavefunction*¹⁴,

$$\phi(\mathbf{x}, t) \equiv N^{1/2} \varphi(\mathbf{x}, t), \quad (2.65)$$

so that $|\phi|^2$ is the particle density, and use the approximation $(N-1)/N \approx 1$. The action is then given by

$$S[\phi] = \int_{t_0}^{t_1} dt \int d\mathbf{x} \phi^*(\mathbf{x}, t) [i\hbar\partial_t - H_{\text{sp}}] \phi(\mathbf{x}, t) - \frac{U_0}{2} |\phi(\mathbf{x}, t)|^4. \quad (2.66)$$

As a final step we apply the variational principle $\delta S / \delta \phi^* = 0$. Dropping the space and time indices for brevity, the functional derivative is

$$\frac{\delta S}{\delta \phi^*} = [i\hbar\partial_t - H_{\text{sp}}] \phi - U_0 |\phi|^2 \phi; \quad (2.67)$$

we again direct the reader to appendix A.2 for details on computing functional derivatives. Setting this to zero and using the usual form for the single-particle Hamiltonian as in Eq. (2.41) yields the Gross-Pitaevskii equation,

$$i\hbar\partial_t \phi = -\frac{\hbar^2}{2m} \nabla^2 \phi + V\phi + U_0 |\phi|^2 \phi. \quad (2.68)$$

2.3.3 Validity of the GPE

The assumptions underlying the GPE are physically reasonable at very low energies¹⁵ and when the scattering length a_s is much less than the interparticle spacing. This is borne out in the many successful applications to ultracold gas experiments, where the GPE is an important theoretical tool for understanding the dynamics of nearly pure condensates [96].

Nevertheless, we emphasise that the derivation given above is not particularly rigorous from a mathematical point of view: We have said nothing about the size of the time dependent errors incurred by making the product state ansatz of Eq. (2.51); neither have we shown that

¹⁴Note that the condensate wavefunction is not a permissible single-particle wavefunction due to the normalisation convention.

¹⁵It is often said that the GPE is valid at “zero temperature”. Strictly speaking, the GPE does not describe a state with a well-defined temperature because it is time dependent, but we may start from a true zero temperature state and apply a coherent excitation.

the product state is a good initial state for the evolution. Perhaps more severely, the effect of the replacement $V_{\text{int}}(\mathbf{x}) = \frac{U_0}{2}\delta(\mathbf{x})$ has also been glossed over, effectively neglecting the short range correlations [39] that arise from the real potential V_{int} . Significant mathematical effort has been put into resolving these problems and convergence proofs are now available [39, 99] for the limit $N \rightarrow \infty$ with $Na_s = \text{fixed}$, known as the Gross-Pitaevskii or mean-field limit.

It is important to understand the nature of convergence in such proofs. In particular, the ansatz Φ does *not* converge to the exact solution Ψ of the N particle Schrödinger equation in the limit of large N ,

$$\|\Psi(\mathbf{x}, t) - \Phi(\mathbf{x}, t)\| \not\rightarrow 0 \quad \text{as } N \rightarrow \infty. \quad (2.69)$$

Instead, it is convergence of the one body reduced density matrix¹⁶ $\rho(\mathbf{x}, \mathbf{x}') \equiv \langle \Psi | \hat{\psi}^\dagger(\mathbf{x}) \hat{\psi}(\mathbf{x}') | \Psi \rangle$ to $\varphi^*(\mathbf{x})\varphi(\mathbf{x}')$:

$$\left\| \frac{1}{N} \rho(\mathbf{x}, \mathbf{x}') - \varphi^*(\mathbf{x})\varphi(\mathbf{x}') \right\| \rightarrow 0 \quad \text{as } N \rightarrow \infty \quad (2.70)$$

where $\|\cdot\|$ is the trace norm¹⁷. In terms of the Penrose-Onsager criterion for condensation [95], this means that the system converges to a pure condensate in the mean-field limit.

To get a feeling for the importance of using the density matrix rather than the state when proving convergence, consider the N particle state

$$\Phi' \equiv \mathcal{S} \left\{ \left[\prod_{k=1}^{N-1} \varphi(\mathbf{x}_k, t) \right] \varphi_\perp(\mathbf{x}_N, t) \right\}, \quad (2.71)$$

where φ_\perp satisfies $\langle \varphi_\perp | \varphi \rangle = 0$ and \mathcal{S} ensures correct Bose symmetrisation of the wavefunction. Intuitively this state has “nearly all particles in state φ ” and yet it is orthogonal to the product $\Phi = \prod_{k=1}^N \varphi(\mathbf{x}_k, t)$ so that $\|\Phi - \Phi'\|$ is never small, even as $N \rightarrow \infty$. On the other hand, we may show¹⁸ that the reduced density operator for Φ' is

$$\langle \Phi' | \hat{\psi}^\dagger(\mathbf{x}) \hat{\psi}(\mathbf{x}') | \Phi' \rangle = (N-1) \varphi^*(\mathbf{x})\varphi(\mathbf{x}') + \varphi_\perp^*(\mathbf{x})\varphi_\perp(\mathbf{x}') \quad (2.72)$$

so that $\frac{1}{N} \rho(\mathbf{x}, \mathbf{x}')$ clearly converges to $\varphi^*(\mathbf{x})\varphi(\mathbf{x}')$ as $N \rightarrow \infty$.

¹⁶For a well-defined number of atoms N in the pure quantum state Ψ , this definition is equivalent to the definition $\rho(\mathbf{x}, \mathbf{x}') = N \int \cdots \int d\mathbf{x}_2 \cdots d\mathbf{x}_N \Psi^*(\mathbf{x}, \mathbf{x}_2, \dots, \mathbf{x}_N) \Psi(\mathbf{x}', \mathbf{x}_2, \dots, \mathbf{x}_N)$, so that the normalisation is $\int d\mathbf{x} \rho(\mathbf{x}, \mathbf{x}) = N$ [96, §13.5]. This normalisation is conventional in ultracold atoms research so we use it here even though normalisation to the identity is more convenient for the discussion of convergence.

¹⁷The trace norm is defined by $\|A\| = \text{Tr} [(A^\dagger A)^{1/2}]$.

¹⁸In second-quantised notation, $|\Phi'\rangle = |N-1, 1, 0, \dots\rangle$ in the basis $\{\varphi, \varphi_\perp, \dots\}$, so we have

$$\hat{\psi}(\mathbf{x})|\Phi'\rangle = (N-1)^{1/2} \varphi(\mathbf{x})|N-2, 1, 0, \dots\rangle + \varphi_\perp(\mathbf{x})|N-1, 0, \dots\rangle,$$

which allows ρ to be computed immediately. This is a small example of the great practical benefit of working in the second-quantised formalism over having to deal with explicitly symmetrised wavefunctions as in Eq. (2.71).

2.4 The Projected Gross-Pitaevskii equation

The full quantum field evolution given in Eq. (2.50) contains answers to any question we might ask about cold Bose gases: Given a state $|\Psi\rangle$ and a *solution* $\hat{\psi}(\mathbf{x}, t)$ to the field equation, any observable $\hat{O}(t)$ can be written as a function $\hat{O}(\hat{\psi}^\dagger(\mathbf{x}, t), \hat{\psi}(\mathbf{x}, t))$ of the field operator; the expectation value of \hat{O} at some time t is

$$\langle \hat{O}(t) \rangle \equiv \langle \Psi | \hat{O}(\hat{\psi}^\dagger(\mathbf{x}, t), \hat{\psi}(\mathbf{x}, t)) | \Psi \rangle. \quad (2.73)$$

Unfortunately solving the operator equations of motion for $\hat{\psi}$ directly is intractable in all but the simplest of cases. As a consequence there is an extensive literature on methods for approximating the equations and/or using an alternative but more tractable form. Exact reformulations typically recast the problem of computing expectation values as a problem of stochastic sampling from some complicated probability distribution; this is the approach taken by the various flavours of quantum Monte Carlo and the phase space methods commonly used in quantum optics. Approximation methods are many and varied, but usually draw on physical insight about the kinds of quantum states which are of interest to reduce the complexity of a solution. In the following section we describe the classical field approximation in the particular form of the projected Gross-Pitaevskii equation (PGPE), originally developed in Refs. [30, 26].

2.4.1 Conceptual introduction

Under certain conditions a system governed by a quantum field theory can instead be well approximated using a classical field. This depends on both the state of the system and the kinds of measurements we wish to make. The prototypical example is the electromagnetic field: while the underlying theory is quantum mechanical, the classical theory — in the form of Maxwell's equations for the classical electric and magnetic fields — describes a huge range of wave phenomena very successfully.

This begs the question: which states and types of measurements can be adequately described using a classical field? Qualitatively, we can say that classical fields describe only the collective wave-like behaviour of systems that are in reality made up of quanta, and that we should have an appropriately large number of those quanta so that the system is approximately continuous rather than discrete. More formally, this is a difficult question to answer because taking the classical limit of a quantum theory is a subtle business. The standard quantum states that lead to classical wave-like behaviour are the Glauber coherent states [52]¹⁹. Unfortunately, there

¹⁹A single-mode coherent state (conventionally written $|\alpha\rangle$) is defined to be an eigenvector of the annihilation operator, such that $\hat{a}|\alpha\rangle = \alpha|\alpha\rangle$ where α is a complex number. In a similar way, there is a full multimode

are conceptual and practical reasons that make it hard to accept coherent states as a good representation of the actual state found in a cold gas experiment. Conceptually, a coherent state is a superposition of states with different numbers of quanta, which has led to an ongoing debate about whether such states can even be generated in principle²⁰ [7]. From a practical point of view, the number statistics of a real system is likely to include additional classical noise which would be inconsistent with the Poissonian number statistics of a coherent state [62]. These objections can be avoided if we treat the state as a classical ensemble of coherent states, as discussed in more detail at the end of the next section.

The electromagnetic field is the most obvious example of a classical field arising out of an underlying quantum theory, but the same kind of approximation can be made for field theories describing ultracold gases of atoms: While the details of the equations of motion and the Hamiltonian differ, both theories ultimately describe the same kinds of objects — indistinguishable bosonic quanta having both wave-like and particle-like properties. For an excellent discussion of the nature of quanta, we direct the reader to Ref. [122].

One might worry that we now have applied the label “classical” to two mutually contradictory theories, both ostensibly describing massive particles: On the one hand, we have classical Newtonian particle mechanics, and on the other a classical field theory. We note that this is nothing other than a manifestation of wave particle duality and the fact that different classical limits are applicable to different situations [101].

To summarise, gases of cold bosons are not well described by the familiar limit of classical *particle* mechanics. As a result, we are led to the quantum theory of many particles in the form of a quantum field theory. Being difficult to solve, this theory is approximated — but not by the classical theory that we had to begin with. Instead the appropriate limit is a classical *field* theory emphasising the collective wave-like aspects of the atomic ensemble.

2.4.2 Derivation of the PGPE

Having said a few general things about the use of classical field theories to describe systems of massive particles, we turn to describing the PGPE that will be used in chapter 3 of this thesis. We stated earlier that classical fields are a good approximation when the number of particles is large. More precisely, it is the number of particles per mode $\langle n_i \rangle$ which is of importance because the errors are of order $1/\langle n_i \rangle$ as argued in Ref. [71]. As a result, we want each mode to be highly occupied, that is, $\langle n_i \rangle \gg 1$ for all i . A system obeying this condition is necessarily

coherent state $|\Psi\rangle$ for each complex-valued field $\Psi(\mathbf{x}, t)$ such that $\hat{\psi}|\Psi\rangle = \Psi(\mathbf{x}, t)|\Psi\rangle$.

²⁰In fact, the possibility of creating coherent states has been called into doubt even in optics [87] where the photon number is not conserved. We direct the reader to Ref. [7] for a presentation of both sides of the debate about the reality of coherent states, and a possible resolution.

highly Bose degenerate (see the start of chapter 1). We note that the number of particles per mode is the same as the density of particles in phase space when the modes in question are the momentum eigenstates — that is, for a homogeneous system. For this reason the terms Bose degeneracy and phase space density are often used interchangeably.

In a spatially finite system we have an infinite number of field modes, so only a finite number of those modes can be highly Bose degenerate as required for the classical field approximation. To deal with this issue, the PGPE formalism splits the set of modes into two subsets, now conventionally labelled the **C** and **I** regions [14]. The **C** or *c-field* region²¹ is chosen to contain the highly occupied modes, and is simulated using a classical field. The **I** or *incoherent* region contains the remaining modes and is assumed to be thermalised. In thermal equilibrium the occupation depends on the mode energy, so the split between **C** and **I** regions is conveniently implemented using an appropriate energy cutoff ϵ_{cut} to be defined more precisely later. We emphasise that the choice of splitting depends on the total number of atoms, so any given splitting is specialised for a subset of the possible quantum states.

There are two well known methods for deriving the PGPE. The first of these is a heuristic approach using what we would like to call “dequantisation” — the assertion that one may simply replace $\hat{\psi}(\mathbf{x})$ by a classical field $\psi_{\mathbf{C}}(\mathbf{x})$ in the operator equations of motion. This is the approach taken, for example, in Ref. [30] and will be discussed further below due to its intuitive appeal.

The second approach uses the truncated Wigner function formalism to derive the PGPE as originally described in Ref. [47] and reviewed in depth in Ref. [14]. The Wigner function method is attractive because it puts the classical field approximation on firmer mathematical ground and allows for systematic treatment of additional quantum behaviour as well as interactions with the **I** region. Nevertheless, we will not need these additional features and there is a significant cost in mathematical machinery, so we will not delve into the details here. We instead direct the reader to Ref. [14, §2] for a clear and detailed explanation.

For a given partitioning of the modes into sets **C** and **I**, we define a pair of projection operators onto the subspaces spanned by the sets of mode functions $\{\phi_i : i \in \mathbf{C}\}$ and $\{\phi_i : i \in \mathbf{I}\}$, respectively:

$$\mathcal{P}_{\mathbf{C}} \{f\}(\mathbf{x}) = \sum_{i \in \mathbf{C}} \phi_i(\mathbf{x}) \int d\mathbf{x}' \phi_i^*(\mathbf{x}') f(\mathbf{x}'), \quad (2.74)$$

$$\mathcal{P}_{\mathbf{I}} \{f\}(\mathbf{x}) = \sum_{i \in \mathbf{I}} \phi_i(\mathbf{x}) \int d\mathbf{x}' \phi_i^*(\mathbf{x}') f(\mathbf{x}'). \quad (2.75)$$

Note that the projectors act on *spatial* functions $f(\mathbf{x})$, as distinct from the creation and anni-

²¹The **C** region has also been called the *coherent* region in earlier papers.

hilation operators that act on the Fock space. The field $\hat{\psi}$ may then be split according to

$$\hat{\psi}(\mathbf{x}, t) = \hat{\psi}_{\mathbf{C}}(\mathbf{x}, t) + \hat{\psi}_{\mathbf{I}}(\mathbf{x}, t), \quad (2.76)$$

where $\hat{\psi}_{\mathbf{C}} \equiv \mathcal{P}_{\mathbf{C}}\{\hat{\psi}\}$ and $\hat{\psi}_{\mathbf{I}} \equiv \mathcal{P}_{\mathbf{I}}\{\hat{\psi}\}$. At this stage no approximation has been made, and one can easily derive a coupled pair of equations of motion for $\hat{\psi}_{\mathbf{C}}$ and $\hat{\psi}_{\mathbf{I}}$ that taken together are equivalent to the full evolution. Applying the projector $\mathcal{P}_{\mathbf{C}}$ to the field equation (2.50) gives

$$\begin{aligned} i\hbar \frac{\partial \hat{\psi}_{\mathbf{C}}}{\partial t} = & H_{\text{sp}} \hat{\psi}_{\mathbf{C}} + U_0 \mathcal{P}_{\mathbf{C}} \left\{ \hat{\psi}_{\mathbf{C}}^\dagger \hat{\psi}_{\mathbf{C}} \hat{\psi}_{\mathbf{C}} + \right. \\ & \hat{\psi}_{\mathbf{I}}^\dagger \hat{\psi}_{\mathbf{C}} \hat{\psi}_{\mathbf{C}} + 2\hat{\psi}_{\mathbf{C}}^\dagger \hat{\psi}_{\mathbf{I}} \hat{\psi}_{\mathbf{C}} + \\ & \hat{\psi}_{\mathbf{C}}^\dagger \hat{\psi}_{\mathbf{I}} \hat{\psi}_{\mathbf{I}} + 2\hat{\psi}_{\mathbf{I}}^\dagger \hat{\psi}_{\mathbf{I}} \hat{\psi}_{\mathbf{C}} + \\ & \left. \hat{\psi}_{\mathbf{I}}^\dagger \hat{\psi}_{\mathbf{I}} \hat{\psi}_{\mathbf{I}} \right\} \end{aligned} \quad (2.77)$$

where it is assumed that the basis ϕ_i is an eigenbasis of the single-particle Hamiltonian so that $\mathcal{P}_{\mathbf{C}}$ commutes with H_{sp} giving $\mathcal{P}_{\mathbf{C}}\{H_{\text{sp}}\hat{\psi}\} = H_{\text{sp}}\mathcal{P}_{\mathbf{C}}\{\hat{\psi}\} = H_{\text{sp}}\hat{\psi}_{\mathbf{C}}$.

We now proceed by the rather gross approximation of discarding the terms coupling the \mathbf{I} region to the evolution of $\hat{\psi}_{\mathbf{C}}$ to obtain

$$i\hbar \frac{\partial \hat{\psi}_{\mathbf{C}}}{\partial t} = H_{\text{sp}} \hat{\psi}_{\mathbf{C}} + U_0 \mathcal{P}_{\mathbf{C}} \left\{ \hat{\psi}_{\mathbf{C}}^\dagger \hat{\psi}_{\mathbf{C}} \hat{\psi}_{\mathbf{C}} \right\}. \quad (2.78)$$

This has the great advantage of providing a closed system which makes numerical work much simpler, but is hard to justify in general (see [26] for further comments). Nevertheless, numerical experience shows that the isolated \mathbf{C} region as described by the PGPE evolves to thermal equilibrium [30], and a range of useful results have been obtained (see [14, §3.3–4] for several examples). We note that the step of discarding all interactions with the \mathbf{I} region is a defining feature of the PGPE regardless of which derivation is used. Retaining these terms is necessary in many cases, in particular for realistic simulations of condensate formation. In such cases it is necessary to use a more powerful alternative, a role filled by the so-called stochastic PGPE (SPGPE) [47].

The final step in the derivation is to make the classical field approximation, that is, to take the classical limit of the field. On a purely formal level this is achieved by simply making the replacement $\hat{\psi}_{\mathbf{C}}(\mathbf{x}, t) \rightarrow \psi_{\mathbf{C}}(\mathbf{x}, t)$ in the equations of motion²², where $\psi_{\mathbf{C}}$ is a complex-valued

²²Equivalently, this is a replacement of each mode operator \hat{a}_i with a complex amplitude c_i for all $i \in \mathbf{C}$.

classical field. We then obtain the PGPE,

$$i\hbar \frac{\partial \psi_{\mathbf{C}}}{\partial t} = H_{\text{sp}} \psi_{\mathbf{C}} + U_0 \mathcal{P}_{\mathbf{C}} \left\{ \psi_{\mathbf{C}}^* \psi_{\mathbf{C}} \psi_{\mathbf{C}} \right\}, \quad (2.79)$$

which describes classical evolution of the highly occupied field modes. While this derivation is expedient and intuitive, it leaves much to be desired from a mathematical standpoint.

The most basic physical justification for the replacement $\hat{\psi}_{\mathbf{C}} \rightarrow \psi_{\mathbf{C}}$ seems to be to regard it as a “dequantisation” or the opposite of canonical quantisation²³. The recipe of canonical quantisation produces a quantum theory from a classical one by replacing the classical canonical position and momentum variables with operators that satisfy the usual commutation relations (see, for example, [41, §1.4]). We imagine dequantisation as the opposite process — forming a classical limit by replacing the quantum field with a classical one.

One conceptual problem with this story is that the nature of the quantum and classical fields are completely different. As discussed in Ref. [122, Ch. 5], a classical field theory uses field equations to describe the state of the system but the quantum field is not a state-like object. Instead the quantum field is more akin to a Green’s function: a solution $\hat{\psi}(\mathbf{x}, t)$ encodes all possible evolutions independently of any particular quantum state. In light of this, it is confusing to imagine replacing or somehow approximating the field operator by a classical field. Clearly some further mathematical justification is required.

The classic mathematical procedure for taking the classical limit is to argue that the class of “classical-like” quantum states of interest are well approximated by coherent states [52]. From this viewpoint the replacement $\hat{\psi}_{\mathbf{C}} \rightarrow \psi_{\mathbf{C}}$ is implemented by assuming the state is coherent and taking the expectation value of both sides of Eq. (2.78). For example, on the left hand side we have

$$\left\langle \Psi \left| i\hbar \frac{\partial \hat{\psi}_{\mathbf{C}}(\mathbf{x}, t)}{\partial t} \right| \Psi \right\rangle = i\hbar \frac{\partial \langle \Psi | \hat{\psi}_{\mathbf{C}}(\mathbf{x}, t) | \Psi \rangle}{\partial t} = i\hbar \frac{\partial \psi_{\mathbf{C}}(\mathbf{x}, t) \langle \Psi | \Psi \rangle}{\partial t} = i\hbar \frac{\partial \psi_{\mathbf{C}}(\mathbf{x}, t)}{\partial t}, \quad (2.80)$$

where $\hat{\psi}_{\mathbf{C}}(\mathbf{x}, t) | \Psi \rangle = \psi_{\mathbf{C}}(\mathbf{x}, t) | \Psi \rangle$ because $| \Psi \rangle$ is a coherent state²⁴. This approach is known as the *broken symmetry* or *mean-field* approach since the average of the quantum field $\langle \hat{\psi}_{\mathbf{C}} \rangle = \psi_{\mathbf{C}}$ is nonzero for a coherent state.

The use of coherent states is attractively simple and provides a straightforward way to take the classical limit. On the other hand, the coherent states are not obviously adequate for cold

²³Canonical quantisation has a somewhat unfortunate name: It is not the canonical method of quantisation, but a method of quantisation making use of classical canonical coordinates.

²⁴The standard notation $\hat{\psi}_{\mathbf{C}}$ and $\psi_{\mathbf{C}}$ is somewhat unfortunate here, since it suggests that the state-like object $\psi_{\mathbf{C}}$ arises as an approximation to the quantum field $\hat{\psi}_{\mathbf{C}}$. This is not the case: it is more accurate to view $\psi_{\mathbf{C}}$ as a representation of the state $| \Psi \rangle$. A less confusing notation might be to write $\Psi(\mathbf{x}, t)$ in place of $\psi_{\mathbf{C}}(\mathbf{x}, t)$ so that $\hat{\psi}_{\mathbf{C}} | \Psi \rangle = \Psi(\mathbf{x}, t) | \Psi \rangle$.

atom experiments, for the reasons noted in Sec. 2.4.1. Luckily, we can avoid those particular objections if we generalise the state to a statistical mixture of coherent states as in Ref. [71]. The density matrix is then expressed as

$$\hat{\rho} = \int \mathcal{D}\alpha P(\alpha) |\alpha\rangle\langle\alpha|, \quad (2.81)$$

where $P(\alpha)$ is a distribution²⁵ over the set of classical fields α . With this representation at hand, computing the dynamics involves evolving each coherent state independently according to the classical field approximation [71]. Observables become ensemble averages at time t over the ensemble described by P .

The procedure outlined above is not a particularly satisfying *mathematical* account of the PGPE theory, but it provides a flavour for the kinds of physical arguments that have been used in the field. With that in mind, we again direct the reader to Ref. [14] for a description of the Wigner function formalism that puts the method on more solid theoretical ground: It shows us more clearly which ensemble we should be using, which terms must be neglected in forming the equations of motion, and directions for extending the theory to deal with physical situations where the PGPE is not adequate.

2.4.3 Ergodicity and thermal averages

We are often interested in the properties of the system at a particular temperature — that is, the expectation values of observations of the *thermal ensemble*. To compute such thermal expectation values it is sufficient to sample any one of the standard statistical ensembles at the desired temperature²⁶. The microcanonical ensemble is simplest to deal with in this case — microcanonical averages should be taken over the hypersurface of constant energy, and the PGPE is energy conserving. If the classical energy is given by

$$H_{\mathbf{C}}[\psi_{\mathbf{C}}] = \int d\mathbf{x} \left(\psi_{\mathbf{C}}^* H_{\text{sp}} \psi_{\mathbf{C}} + \frac{U_0}{2} |\psi_{\mathbf{C}}|^4 \right), \quad (2.82)$$

²⁵As written here, P is actually the density in the Glauber-Sudarshan P representation [52], and as a result is much more general than one would guess at first sight. Not all density matrices have a well behaved P representation (in general P can be negative and extremely singular [16]), but this is not a problem for the states of interest since — roughly speaking — the more classical a state is, the better behaved is the associated P distribution. We note that the Wigner function mentioned above is equal to the P distribution smoothed by a Gaussian convolution.

²⁶The various statistical ensembles are equivalent in the thermodynamic limit. It is worth keeping in mind that they have different fluctuation properties for the mesoscopic numbers of particles we are dealing with here.

then microcanonical averages are taken over all $\psi_{\mathbf{C}}$ with respect to the phase space density [123]

$$P[\psi_{\mathbf{C}}; E] = \begin{cases} \text{const.} & \text{where } H_{\mathbf{C}}[\psi_{\mathbf{C}}] = E \text{ and other macroscopic constraints are satisfied} \\ 0 & \text{elsewhere.} \end{cases} \quad (2.83)$$

We mention other macroscopic constraints because there may be additional constants of motion and we want any averages to take these into account. For example, the momentum is conserved in a homogeneous system in which case the ensemble of interest includes only microstates with the system at rest.

The crucial step in computing thermal averages is to assume that the PGPE is *ergodic* [30] which allows us to convert phase space averages into time averages:

$$\langle F \rangle = \int \mathcal{D}\psi_{\mathbf{C}} P[\psi_{\mathbf{C}}; E] F[\psi_{\mathbf{C}}] \quad (2.84)$$

$$= \lim_{\tau \rightarrow \infty} \frac{1}{\tau} \int_0^\tau dt F[\psi_{\mathbf{C}}(t)] \quad (2.85)$$

where F is some functional of the classical field representing an observable. In practice the time average is implemented numerically by sampling the motion at discrete time intervals $\{t_j\}$ and forming the sum

$$\langle F \rangle \approx \frac{1}{M} \sum_{j=1}^M F[\psi_{\mathbf{C}}(t_j)] \quad (2.86)$$

for large M . This recipe provides a convenient and efficient way to sample states from the microcanonical ensemble.

Not all quantities may be easily written as functionals of the field $\psi_{\mathbf{C}}$. In particular we note that derivatives of entropy such as the temperature (T) and chemical potential ($\mu_{\mathbf{C}}$) are calculated by time-averaging appropriate quantities constructed from the Hamiltonian in Eq. (2.82) using the Rugh approach [110]. The detailed implementation of the Rugh formalism for the PGPE is rather technical and we refer the reader to Refs. [29, 27] for additional details of this procedure.

It is instructive to connect the discussion of ergodic averaging to the Wigner function version of the PGPE derivation. For a thermal state the initial Wigner function is very delocalised, in contrast to the near delta function required for a straightforward single-trajectory interpretation. Sampling directly from such a distribution is difficult except in the limit of very low or high temperatures where the Hamiltonian can be approximately diagonalised [14]. Happily, the PGPE with ergodic averaging avoids this problem by relying on the *dynamics* to sample

the distribution correctly.

2.4.4 Treatment of the **I** region

At temperatures near the BEC transition the number of atoms in the **I** region is significant, as observed in chapter 3 (see also [15] for three-dimensional systems). For realistic comparisons with experiment these need to be taken into account.

A simple way to deal with the **I** region atoms is to assume a gas of uncorrelated semiclassical bosons in thermal equilibrium, interacting only indirectly via the average particle density [10, 28]. Under these assumptions the **I** region can be described by an approximate single-particle Wigner function $W_{\mathbf{I}}$. (Note that this is a *single*-particle Wigner function, distinct from the multi-particle Wigner function referred to previously in section 2.4.2.) The appropriate single-particle Wigner function is positive and can be interpreted as a classical probability distribution over phase space:

$$W_{\mathbf{I}}(\mathbf{k}, \mathbf{x}) = \frac{1}{(2\pi)^D} \frac{1}{e^{[E_{\text{HF}}(\mathbf{k}, \mathbf{x}) - \mu]/k_B T} - 1}. \quad (2.87)$$

In this equation D is the number of dimensions, \mathbf{k} is the wavevector and $\mu = \mu_{\mathbf{C}} + 2U_0 n_{\mathbf{I}}$ is the chemical potential. The Hartree-Fock energy in this expression is given by

$$E_{\text{HF}}(\mathbf{k}, \mathbf{x}) = \frac{\hbar^2 \mathbf{k}^2}{2m} + V(\mathbf{x}) + 2U_0 [n_{\mathbf{C}}(\mathbf{x}) + n_{\mathbf{I}}(\mathbf{x})], \quad (2.88)$$

where $n_{\mathbf{C}}$ and $n_{\mathbf{I}}$ are the densities of the **C** and **I** region atoms [10]; the temperature and chemical potential are calculated from the **C** region. When the potential V is not constant the unknown density $n_{\mathbf{I}}$ is spatially varying and must be calculated self-consistently [28] using

$$n_{\mathbf{I}}(\mathbf{x}) = \int_{E_{\text{HF}}(\mathbf{k}, \mathbf{x}) > \epsilon_{\text{cut}}} d\mathbf{k} W_{\mathbf{I}}(\mathbf{k}, \mathbf{x}). \quad (2.89)$$

This complication disappears in the homogeneous case relevant to the work in chapter 3.

With the semiclassical Wigner function at hand, any observable $F(\mathbf{k}, \mathbf{x})$ may be calculated using a phase space average

$$\langle F \rangle_{\mathbf{I}} = \int_{E_{\text{HF}}(\mathbf{k}, \mathbf{x}) > \epsilon_{\text{cut}}} d\mathbf{x} d\mathbf{k} W_{\mathbf{I}}(\mathbf{k}, \mathbf{x}) F(\mathbf{k}, \mathbf{x}). \quad (2.90)$$

Note that the nontrivial region of integration satisfying $E_{\text{HF}}(\mathbf{k}, \mathbf{x}) > \epsilon_{\text{cut}}$ is chosen to avoid counting atoms that have already been taken into account via the **C** region simulation.

2.4.5 Comparison with the GPE

We briefly contrast the PGPE formalism with the GPE, as the relationship between these is a recurring source of confusion: Given such apparently similar equations, why is it claimed that the PGPE describes all modes at temperatures up to the order of the transition temperature, while the GPE describes only the condensate mode at zero temperature? To further emphasise the point of similarity, recall that there are various ways to derive the GPE with subtly different interpretations as to the exact state described. In Sec. (2.3) we presented the GPE as arising from the product state ansatz of Eq. (2.51), but it can just as easily be derived by assuming a time dependent coherent state²⁷. This second approach is confusingly similar to the “dequantisation” step in our derivation of the PGPE when interpreted in terms of coherent states.

The real heart of the difference between these formalisms is the ensemble average that must be taken when using the PGPE to describe finite temperature states of the field. While this is implemented as an ergodic average, certain non-thermal states could also in principle be simulated by considering a full set of trajectories drawn from an initial non-thermal distribution. Averages would then be taken over this set of trajectories at some time²⁸. At finite temperature there must be some portion of non-condensed atoms, and indeed the PGPE is able to simulate these: The ensemble average results in a single particle density matrix with a largest eigenvalue significantly less than one; this is a non-pure condensate in the Penrose-Onsager sense.

In contrast, the GPE simulates a single trajectory representing the motion of the condensate: The single particle density matrix becomes arbitrarily close to pure in the GP limit, as made precise in Sec. 2.3.3. The types of states that the GPE simulates are thus small perturbations of the zero temperature stationary solution²⁹.

While the projection operator is required to provide control over the set of modes to be approximated, the arguments above show that it is not at the heart of the conceptual differences between the GPE and PGPE. Nevertheless, we emphasise that a projector is essential to clearly

²⁷The sense in which the GPE approximates the true state of the system is addressed in Sec. 2.3.3. We point out here that it is consistent for both the coherent *and* product state ansätze to approximate the true many particle state in the relatively weak sense of Eq. (2.70).

Furthermore, distinguishing between these apparently very different states in an experiment can be surprisingly difficult. This is nicely demonstrated in Ref. [125] that analyses an interference experiment, finding that the difference in interference fringe contrast is of order $1/N$, and very difficult to measure for even moderate numbers of particles N . This is reminiscent of the approximation $N/(1+N) \approx 1$ that must be made in deriving the GPE via the product state ansatz.

²⁸Indeed, this is the approach generally taken by the whole group of powerful methods arising from the truncated Wigner function formalism.

²⁹Non-stationary solutions are not, strictly speaking, at zero temperature since simple definitions of thermal equilibrium imply a stationary state. A dynamical definition of temperature such as given in Ref. [110] would presumably assign a small but nonzero temperature to such states.

define the set of modes in the **I** region. Without careful treatment of these high energy modes, one can only hope to achieve qualitative agreement with experimental observations.

Vortex pairing in two-dimensional Bose gases

In this chapter we investigate finite temperature 2D Bose gases using the PGPE, with a view to understanding the relation between BEC and BKT physics in finite-sized systems. We calculate several physical properties including the amount of vortex pairing, the condensate and superfluid fractions, and the functional form of the spatial correlations. We also relate our simulation to the experimental measurements described in Ref. [60].

3.1 Introduction

As discussed in section 1.2, a two-dimensional homogeneous Bose gas does not undergo the BEC transition. Nevertheless, 2D Bose gases do display superfluid behaviour in the presence of interactions, due to a vortex pairing phase transition known as the Berezinskii-Kosterlitz-Thouless (BKT) transition. Evidence for the BKT transition has been found in experimental realisations of the 2D Bose gas in several studies [121, 60, 77, 117, 23]. We make particular note of the experiment described in Ref. [60], which was carried out at ENS in Paris, and which we will refer to when choosing parameters for our study.

Experiments in the 2D regime present a new challenge for theory as strong fluctuations invalidate mean-field theories (see, for example, [105, 106, 104, 12, 83, 113, 50, 116]), and only recently have quantum Monte Carlo [68, 67] and classical field (c-field) [118, 119, 13] methods been developed that are directly applicable to the experimental regime.

In the current chapter we study a uniform Bose gas of finite spatial extent and parameters corresponding to current experiments. To analyse this system we use the PGPE, which is well suited to studying finite temperature Bose fields with many highly occupied modes. We examine two important applications: First, we provide a quantitative validation of the interference technique used in the ENS experiment to determine the nature of two-point correlation in the system. To do this we simulate the interference pattern generated by allowing two independent 2D systems to expand and interfere. Applying the experimental fitting procedure to analyse the interference pattern, we can extract the inferred two-point correlations which we then compare

against the *in situ* correlations that we calculate directly. Second, we examine the correlations between vortices and antivortices in the system to directly quantify the emergence of vortex-antivortex pairing in the low temperature phase. A similar study was made by Giorgetti *et al.* using a semiclassical field technique [51]. We find results for vortex number and vortex pair distributions consistent with their results, and we show how a coarse-graining procedure can be used to reveal the unpaired vortices in the system.

3.2 Formalism

Here we consider a dilute 2D Bose gas described by the Hamiltonian

$$\hat{H} = \int d^2\mathbf{x} \hat{\psi}^\dagger \left\{ -\frac{\hbar^2 \nabla_{\mathbf{x}}^2}{2m} \right\} \hat{\psi} + \frac{\hbar^2 g}{2m} \int d^2\mathbf{x} \hat{\psi}^\dagger \hat{\psi}^\dagger \hat{\psi} \hat{\psi}, \quad (3.1)$$

where $\mathbf{x} = (x, y)$. This is simply Eq. (2.47), but specialised to the two-dimensional homogeneous case. We take the two-dimensional geometry to be realised by tight confinement in the z direction that restricts atomic occupation to the lowest z mode. The dimensionless 2D coupling constant is

$$g = \frac{\sqrt{8\pi}a}{a_z}, \quad (3.2)$$

with a_z the spatial extent of the z mode¹ and a the s-wave scattering length. We will assume that $a_z \gg a$ so that the scattering is approximately three-dimensional [97], a condition well-satisfied in the ENS and NIST experiments [121, 60, 77, 23]. For reference, the ENS experiment reported in Ref. [60] had $g \approx 0.15$, whereas in the NIST experiments $g \approx 0.02$ [23].

In contrast to experiments we focus here on the uniform case; no trapping potential in the xy -plane is considered. We perform finite-sized calculations corresponding to a square system of size L with periodic boundary conditions. Working in the finite-size regime simplifies the simulations and is more representative of current experiments. We note that the thermodynamic limit corresponds to taking $L \rightarrow \infty$ while keeping the density, $n = \langle \hat{\psi}^\dagger \hat{\psi} \rangle$, constant.

3.2.1 Review of BKT physics

The BKT superfluid phase has several distinctive characteristics, which we briefly review.

¹For example, for tight harmonic confinement of frequency ω_z we have $a_z = \sqrt{\hbar/m\omega_z}$.

First-order correlations

Below the BKT transition the first-order correlations decay according to an inverse power law:

$$g^{(1)}(\mathbf{x}, \mathbf{x}') \propto \|\mathbf{x} - \mathbf{x}'\|^{-\alpha}. \quad (3.3)$$

Systems displaying such *algebraic decay* are said to exhibit *quasi-long-range order* [19]. This is in contrast to both the high temperature (disordered phase) in which the correlations decay exponentially, and long-range ordered case of the 3D Bose gas in which $g^{(1)} \rightarrow \text{const.}$ for $\|\mathbf{x} - \mathbf{x}'\| \rightarrow \infty$.

Superfluid density

Nelson and Kosterlitz [92] found that the exponent of the algebraic decay is related to the ratio of the superfluid density and temperature. To within logarithmic corrections

$$\alpha(T) = \frac{1}{\lambda_{\text{dB}}^2 \rho_s(T)}, \quad (3.4)$$

where ρ_s is the superfluid density and λ_{dB} is the thermal de Broglie wavelength (Eq. (1.1)). Furthermore, Nelson and Kosterlitz showed that this ratio converges to a universal constant as the transition temperature, T_{KT} , is approached from below: $\lim_{T \rightarrow T_{\text{KT}}^-} \alpha(T) = 1/4$ (i.e., $\rho_s \lambda_{\text{dB}}^2 = 4$). Thus, the superfluid fraction undergoes a universal jump from $\rho_s(T_{\text{KT}}^+) = 0$ to $\rho_s(T_{\text{KT}}^-) = 4/\lambda_{\text{dB}}^2$ as the temperature decreases through T_{KT} .

Vortex binding transition

Another important indicator of the BKT transition is the behaviour of topological excitations, which are quantised vortices and antivortices in the case of a Bose gas. A single vortex has energy that scales with the logarithm of the system size. At low temperatures this means that the free energy for a single vortex is infinite (in the thermodynamic limit), and vortices cannot exist in isolation. As originally argued in Ref. [75], the entropic contribution to the free energy also scales logarithmically with the system size, and will dominate the free energy at high temperatures allowing unbound vortices to proliferate. This argument provides a simple estimate for the BKT transition temperature.

Although unbound vortices are thermodynamically unfavoured at $T < T_{\text{KT}}$, bound pairs of counter-rotating vortices may exist because the total energy of such a pair is finite². This leads to a distinctive qualitative characterisation of the BKT transition: as the temperature

²The vortex-antivortex pair energy depends on the pair size rather than the system size.

increases through T_{KT} pairs of vortices unbind.

Location of the BKT transition in the dilute Bose gas

While the relation $\rho_s(T_{\text{KT}}^-) = 4/\lambda_{\text{dB}}^2$ between the superfluid density and temperature at the transition is universal, the total density, n , at the transition is not. General arguments [103, 72, 43] suggest that the transition point for the dilute uniform 2D Bose gas is given by

$$(n\lambda_{\text{dB}}^2)_{\text{KT}} = \ln\left(\frac{\xi}{g}\right), \quad (3.5)$$

where ξ is a constant. Prokofév, Ruebenacker and Svistunov [105, 106] studied the homogeneous Bose gas using Monte Carlo simulations of an equivalent classical ϕ^4 model on a lattice. Using an extrapolation to the infinite-sized system, they computed a value for the dimensionless constant, $\xi = 380 \pm 3$. By inverting Eq. (3.5), we obtain the BKT critical temperature for the infinite system

$$T_{\text{KT}}^\infty = \frac{2\pi\hbar^2 n}{mk_B \ln(\xi\hbar^2/mg)}. \quad (3.6)$$

We use the superscript ∞ to indicate that this result holds in the thermodynamic limit.

3.3 Method

3.3.1 c-field and incoherent regions

We briefly outline some specifics regarding how the PGPE formalism described in section 2.4 is applied to the two-dimensional homogeneous problem. In the homogeneous case, the fields ψ_{C} and $\hat{\psi}_{\text{I}}$ are defined as the low and high energy projections of the full quantum field operator, separated by the cutoff wave vector K . In our theory this cutoff is implemented in terms of the plane wave eigenstates $\{\varphi_{\mathbf{n}}(\mathbf{x})\}$ of the time-independent single-particle Hamiltonian, that is,

$$\varphi_{\mathbf{n}}(\mathbf{x}) = \frac{1}{L} e^{-i\mathbf{k}_{\mathbf{n}} \cdot \mathbf{x}}, \quad (3.7)$$

$$\mathbf{k}_{\mathbf{n}} = \frac{\pi}{L} \mathbf{n}, \quad (3.8)$$

with $\mathbf{n} = (n_x, n_y) \in \mathbb{Z}^2$. The fields are thus defined by

$$\psi_{\text{C}}(\mathbf{x}) \equiv \sum_{\mathbf{n} \in \text{C}} c_{\mathbf{n}} \varphi_{\mathbf{n}}(\mathbf{x}), \quad (3.9)$$

$$\hat{\psi}_{\text{I}}(\mathbf{x}) \equiv \sum_{\mathbf{n} \in \text{I}} \hat{a}_{\mathbf{n}} \varphi_{\mathbf{n}}(\mathbf{x}), \quad (3.10)$$

where the $\hat{a}_{\mathbf{n}}$ are Bose annihilation operators, the $c_{\mathbf{n}}$ are complex amplitudes, and the sets of quantum numbers defining the regions are

$$\mathbf{C} = \{\mathbf{n}: \|\mathbf{k}_{\mathbf{n}}\| \leq K\}, \quad (3.11)$$

$$\mathbf{I} = \{\mathbf{n}: \|\mathbf{k}_{\mathbf{n}}\| > K\}. \quad (3.12)$$

Choice of \mathbf{C} region

In general, the applicability of the PGPE approach to describing the finite temperature gas relies on an appropriate choice for K , so that the modes at the cutoff have an average occupation of order unity. In this work we choose an average of five or more atoms per mode using a procedure discussed in appendix B.1. This choice means that all the modes in \mathbf{C} are appreciably occupied, justifying the classical field replacement $\hat{a}_{\mathbf{n}} \rightarrow c_{\mathbf{n}}$. In contrast the \mathbf{I} region contains many sparsely occupied modes that are particle-like and would be poorly described using a classical field approximation. Because our 2D system is critical over a wide temperature range, additional care is needed in choosing \mathbf{C} . Typically strong fluctuations occur in the infrared modes up to the energy scale $\hbar^2 g n / m$. Above this energy scale the modes are well described by mean-field theory (see, for example, the discussion in [74, 105]). For the results we present here, we have

$$\frac{\hbar^2 K^2}{2m} \gtrsim \frac{\hbar^2 g}{m} n \quad (3.13)$$

for simulations around the transition region and at high temperature. At temperatures well below T_{KT} , the requirement of large modal occupation near the cutoff competes with this condition and we favour the former at the expense of violating Eq. (3.13).

PGPE treatment of \mathbf{C} region

Specialising the PGPE (Eq. (2.79)) to 2D we have the equation of motion for $\psi_{\mathbf{C}}$

$$i\hbar \frac{\partial \psi_{\mathbf{C}}}{\partial t} = -\frac{\hbar^2 \nabla_{\mathbf{x}}^2}{2m} \psi_{\mathbf{C}} + \frac{\hbar^2 g}{m} \mathcal{P}_{\mathbf{C}} \{ |\psi_{\mathbf{C}}|^2 \psi_{\mathbf{C}} \}, \quad (3.14)$$

where the projection operator

$$\mathcal{P}_{\mathbf{C}} \{ F(\mathbf{x}) \} \equiv \sum_{\mathbf{n} \in \mathbf{C}} \varphi_{\mathbf{n}}(\mathbf{x}) \int d^2 \mathbf{x}' \varphi_{\mathbf{n}}^*(\mathbf{x}') F(\mathbf{x}'), \quad (3.15)$$

formalises our basis set restriction of $\psi_{\mathbf{C}}$ to the \mathbf{C} region. The main approximation used to arrive at the PGPE is to neglect dynamical couplings to the incoherent region [26].

We assume that the evolution under Eq. (3.14) is ergodic [30], so that the microstates $\{\psi_{\mathbf{C}}\}$

generated through time evolution form an unbiased sample of the equilibrium microstates. Time-averaging can then be used to obtain macroscopic equilibrium properties. We generate the time evolution by solving the PGPE with three adjustable parameters: (i) the cutoff wave vector, K , that defines the division between **C** and **I**, and hence the number of modes in the **C** region; (ii) the number of **C** region atoms, $N_{\mathbf{C}}$; (iii) the total energy of the **C** region, $E_{\mathbf{C}}$. The last two quantities, defined as

$$E_{\mathbf{C}} = \int d^2\mathbf{x} \psi_{\mathbf{C}}^* \left(-\frac{\hbar^2 \nabla_{\mathbf{x}}^2}{2m} + \frac{\hbar^2 g}{2m} |\psi_{\mathbf{C}}|^2 \right) \psi_{\mathbf{C}}, \quad (3.16)$$

$$N_{\mathbf{C}} = \int d^2\mathbf{x} |\psi_{\mathbf{C}}(\mathbf{x})|^2, \quad (3.17)$$

are important because they represent constants of motion of the PGPE (Eq. (3.14)), and thus control the thermodynamic equilibrium state of the system.

Obtaining equilibrium properties for the **C** region

To characterise the equilibrium state in the **C** region it is necessary to determine the average density, temperature and chemical potential, which in turn allow us to characterise the **I** region (see section 3.3.2). These and other **C** region quantities can be computed by time-averaging as described in section 2.4.3. For example, the average **C** region density is given by

$$n_{\mathbf{C}}(\mathbf{x}) \approx \frac{1}{M_s} \sum_{j=1}^{M_s} |\psi_{\mathbf{C}}(\mathbf{x}, t_j)|^2, \quad (3.18)$$

where $\{t_j\}$ is a set of M_s times (after the system has been allowed to relax to equilibrium) at which the field is sampled. We typically use 2000 samples from our simulation to perform such averages over a time of ~ 16 s. Another quantity of interest here is the first-order correlation function, which we calculate directly via the expression

$$G_{\mathbf{C}}^{(1)}(\mathbf{x}, \mathbf{x}') \approx \frac{1}{M_s} \sum_{j=1}^{M_s} \psi_{\mathbf{C}}^*(\mathbf{x}, t_j) \psi_{\mathbf{C}}(\mathbf{x}', t_j). \quad (3.19)$$

The temperature (T) and chemical potential ($\mu_{\mathbf{C}}$) are computed using the Rugh approach [110] which was briefly touched upon in section 2.4.3.

A major extension to the formalism of the PGPE made in this thesis is the development of a method for extracting the superfluid fraction, ρ_s , from these calculations. For this we use linear response theory to relate the superfluid fraction to the long wavelength limit of the second order momentum density correlations. An extensive discussion of this approach, and

the numerical methods used to implement it, are presented in chapter 4.

3.3.2 Mean-field treatment of **I** region

Occupation of the **I** region modes, $N_{\mathbf{I}}$, accounts for about 25% of the total number of atoms at temperatures near the phase transition. We assume a time-independent state for the **I** region atoms defined by a Wigner function [91], allowing us to calculate quantities of interest by integrating over the above-cutoff momenta, $k > K$ [10, 28].

Our assumed Wigner function corresponds to the self-consistent Hartree-Fock theory as applied in Ref. [28]. Specialising Eq. (2.87) to two dimensions, this is

$$W_{\mathbf{I}}(\mathbf{k}, \mathbf{x}) = \frac{1}{(2\pi)^2} \frac{1}{e^{(E_{\text{HF}}(\mathbf{k}) - \mu)/k_B T} - 1}, \quad (3.20)$$

where

$$E_{\text{HF}}(\mathbf{k}) = \frac{\hbar^2 \mathbf{k}^2}{2m} + \frac{2\hbar^2 g}{m} (n_{\mathbf{C}} + n_{\mathbf{I}}), \quad (3.21)$$

is the Hartree-Fock energy, $n_{\mathbf{I}}$ is the **I** region density, and $\mu = \mu_{\mathbf{C}} + 2\hbar^2 g n_{\mathbf{I}}/m$ is the chemical potential (shifted by the mean-field interaction with the **I** region atoms). Note that the average densities are constant in the uniform system, so $W_{\mathbf{I}}(\mathbf{k}, \mathbf{x})$ has no explicit \mathbf{x} dependence, however, we include this variable for generality when defining the associated correlation function.

The **I** region density appearing in Eq. (3.21) is given by

$$n_{\mathbf{I}} = \int_{\|\mathbf{k}\| \geq K} d^2 \mathbf{k} W_{\mathbf{I}}(\mathbf{k}, \mathbf{x}), \quad (3.22)$$

with corresponding atom number $N_{\mathbf{I}} = n_{\mathbf{I}} L^2$; total number is simply

$$N = N_{\mathbf{C}} + N_{\mathbf{I}}. \quad (3.23)$$

An analytic expression for $n_{\mathbf{I}}$ and simplified procedure for numerically calculating the first-order correlation function of the **I** region atoms, $G_{\mathbf{I}}^{(1)}$, can be obtained by taking integrals over the phase space. These results are discussed in appendix B.2.

3.3.3 Equilibrium configurations with fixed T and N

Generating equilibrium classical fields with given values of $E_{\mathbf{C}}$ and $N_{\mathbf{C}}$ is straightforward since the PGPE simulates a microcanonical system (see appendix B.1.3). However, we wish to simulate systems with a given temperature and total number. As described in the preceding two sections these can only be determined after a simulation has been performed. In appendix

B.1 we outline a procedure for estimating values of $E_{\mathbf{C}}$ and $N_{\mathbf{C}}$ for desired values of N and T based on a root finding scheme using a Hartree-Fock-Bogoliubov analysis for the initial guess.

3.4 Results

We choose simulation parameters in analogy with the Paris experiment of Hadzibabic *et al.* [60]. This experiment used an elongated atomic cloud of approximately 10^5 ^{87}Rb atoms, with a spatial extent (Thomas-Fermi lengths) of $120\text{ }\mu\text{m}$ and $10\text{ }\mu\text{m}$ along the two loosely trapped x and y directions. The tight confinement in the z direction was provided by an optical lattice.

Although our simulation is for a uniform system, we have chosen similar parameters where possible. Our primary simulations are for a system in a square box with $L = 100\text{ }\mu\text{m}$, with 4×10^5 ^{87}Rb atoms. We also present results for systems with $L = 50\text{ }\mu\text{m}$ and $L = 200\text{ }\mu\text{m}$ at the same density in order to better understand finite-size effects. All simulations are for the case of $g = 0.15$ corresponding to the experimental parameters reported in Ref. [60].

The cutoff wave vector K varied with temperature to ensure appropriate occupation of the highest modes (see section 3.3.1). For the $100\text{ }\mu\text{m}$ system, the number of \mathbf{C} region modes ranged between 559 at low temperatures to 11338 at the highest temperature studied.

3.4.1 Simulation of expanded interference patterns between two systems

In order to make a direct comparison with the experimental results of Ref. [60], we have generated synthetic interference patterns and implemented the experimental analysis technique. Our simulated imaging geometry is identical to that found in Ref. [60], with expansion occurring in the z -direction. The interference pattern is formed in the xz -plane via integration of the density along the y -direction (“absorption imaging”).

Our algorithm for obtaining the interference pattern due to our classical field is very similar to that presented in Ref. [61]. Our above-cutoff thermal cloud is taken into account separately. We consider a pair of fields $\psi_{\mathbf{C}}^{(1)}(x, y), \psi_{\mathbf{C}}^{(2)}(x, y)$ from different times during the simulation, chosen such that the fields can be considered independent. The 3D wavefunction corresponding to each field is reconstructed by assuming a harmonic oscillator ground state in the tight-trapping direction. These two reconstructed fields are spatially separated by $\Delta = 3\text{ }\mu\text{m}$, corresponding to the period of the optical lattice in Ref. [61].

Given this initial state, we neglect atomic interactions and only account for expansion in the tightly-trapped direction. This yields a simple analytical result for the full classical field $\psi_{\mathbf{C}}(x, y, z, \tau)$ at later times. The contribution of the above-cutoff atoms is included by an

incoherent addition of intensities. The result is integrated along the y -direction to simulate the effect of absorption imaging with a laser beam, that is,

$$n_{\text{im}}(x, z) = \int_0^{L'} dy \left[|\psi_{\mathbf{C}}^{(T)}(x, y, z, \tau)|^2 + n_{\mathbf{I}}(x, y, z, \tau) \right], \quad (3.24)$$

$$\psi_{\mathbf{C}}^{(T)} = \psi_{\mathbf{C}}^{(1)}(x, y, z, \tau) + \psi_{\mathbf{C}}^{(2)}(x - \Delta, y, z, \tau). \quad (3.25)$$

Rather than integrate the full field along the y -direction, we use only a slice of length $L' = 10 \mu\text{m}$ in keeping with the experimental geometry of Ref. [60].

The interference patterns, $n_{\text{im}}(x, z)$, generated this way contained fine spatial detail not seen in the experimental images. To make a more useful comparison to experiment it is necessary to account for the finite optical imaging resolution by applying a Gaussian convolution in the xz -plane with standard deviation $3 \mu\text{m}$ [59].

In accordance with the Paris experiment, we use a 22 ms expansion time to generate interference patterns for quantitative analysis (see section 3.4.3). To obtain characteristic interference images for display in Ref. [60], the experiments used a shorter 11 ms expansion [59]. We exhibit examples of interference patterns at various temperatures in Fig. 3.1, for this shorter expansion time. These images show a striking resemblance to the results presented in Ref. [60].

3.4.2 Condensate and superfluid fractions

For a 2D Bose gas in a box we expect a nonzero condensate fraction due to the finite spacing of low-energy modes. A central question is whether we can observe a distinction between the crossover due to Bose condensation and that due to BKT physics. To address this question we have computed both the condensate and superfluid fractions from our dynamical simulations.

The condensate fraction in a homogeneous system is easily identified as the average fractional occupation of the lowest momentum mode. This is directly available from our simulations as a time average of the $\mathbf{k} = \mathbf{0}$ mode of the classical field,

$$f_c = \langle c_{\mathbf{0}}^* c_{\mathbf{0}} \rangle / N. \quad (3.26)$$

Superfluidity is the macroscopic tendency for some fraction f_s of certain fluids to flow without apparent viscosity. The task of connecting this phenomenology to the microscopic theory is not trivial, so extracting the superfluid fraction from dynamical classical field simulations provides a more difficult challenge. Our approach relies on linear response theory to connect the superfluid fraction with the long wavelength limit of the second order momentum density correlations. The details of this technique are presented separately in chapter 4.




Figure 3.1: Synthetic interference patterns generated from the $50\ \mu\text{m}$ grid by simulation of the experimental procedure of Ref. [60]. (a) At low temperatures, $T \approx 0.5T_{\text{KT}}$, the interference fringes are straight. (b) Just below the transition temperature, $T \approx 0.95T_{\text{KT}}$, the fringes become wavy due to decreased spatial phase coherence. Phase dislocations become common at temperatures above the transition, (c) $T \approx 1.05T_{\text{KT}}$, and (d) $T \approx 1.1T_{\text{KT}}$. These “zipper patterns” indicate the presence of free vortices. (e) When simulation of the finite imaging resolution is disabled, the zipper patterns from the field in sub figure (d) are no longer clearly visible; the high frequency details obscure the phase information without providing obvious additional information about the existence of vortex pairs.

Figure 3.2 compares the results for the superfluid and condensate fractions computed on the $100\ \mu\text{m}$ grid. These results are qualitatively similar to the results for the larger and smaller grids. In particular, we note that there is no apparent separation between temperatures at which the superfluid and condensate fractions fall to zero. Also shown in Fig. 3.2 is the condensate fraction for the ideal Bose gas confined to an identical finite-size box in the grand canonical ensemble. The large shift between ideal and computed transition temperatures indicates the effect of interactions in the 2D system. Because the average system density is uniform, this large shift is due to critical fluctuations (also see [74]).

In our calculations we identify the transition temperature, T_{KT} , as where the superfluid fraction falls off most rapidly (i.e., the location of steepest slope on the f_s versus T graph; see Fig. 3.2). As the system size increases, this transition temperature moves toward the value for an infinite-sized system, T_{KT}^∞ [105]. This effect is illustrated by the behaviour of the superfluid fraction in Fig. 3.3.

Also shown in Fig. 3.3 is an alternative calculation of the superfluid fraction based on Eq. (3.4). The two methods are expected to match for temperatures at and slightly below the




Figure 3.2: Condensate fraction (solid dots) and superfluid fraction (crosses) as functions of temperature for the $100 \mu\text{m}^2$ grid. The transition temperature in the thermodynamic limit, T_{KT}^∞ [105], is shown as a vertical dot-dashed line. The vertical dashed line shows our estimate for the transition temperature in the finite system. The thick solid line is the condensate fraction for an ideal Bose gas in the grand canonical ensemble with the same number of atoms and periodic spatial domain. The superfluid fraction becomes negative in places because the extrapolation of the momentum correlations to $\mathbf{k} = 0$ is sensitive to statistical noise at high temperature (see section 4.3 for details).

transition temperature [92].

3.4.3 First-order correlations — algebraic decay

Algebraic decay of the first-order correlations, as described by Eq. (3.3), is a characteristic feature of the BKT phase. Above the BKT transition, the first-order correlations should revert to the exponential decay expected in a disordered phase.

The normalised first-order correlation function, $g^{(1)}$ is defined by

$$g^{(1)}(\mathbf{x}, \mathbf{x}') = \frac{G^{(1)}(\mathbf{x}, \mathbf{x}')}{\sqrt{n(\mathbf{x})n(\mathbf{x}')}}, \quad (3.27)$$

where $G^{(1)}(\mathbf{x}, \mathbf{x}') = \langle \hat{\psi}^\dagger(\mathbf{x}) \hat{\psi}(\mathbf{x}') \rangle$ is the unnormalised first-order correlation function [91]. In a homogeneous isotropic system $g^{(1)}$ depends only the distance $\|\mathbf{x} - \mathbf{x}'\|$ and we may characterise




Figure 3.3: Detail of the superfluid fraction near the transition temperature. Solid dots represent the calculation based on momentum correlations as described in chapter 4. Results for the largest and smallest grids are shown (left and right, respectively). The data for the 100 μm grid is omitted for clarity, but lies between the curves shown as expected. Open circles represent the calculation of the superfluid fraction from the associated fitted values for the decay coefficient α , via Eq. (3.4). The open circles terminate where the power law fitting procedure fails.

the first-order correlations by a function of one variable, $g^{(1)}(x) \equiv g^{(1)}(\|\mathbf{x} - \mathbf{x}'\|) = g^{(1)}(\mathbf{x}, \mathbf{x}')$.

Direct calculation of $g^{(1)}$

In the PGPE formalism the \mathbf{C} and \mathbf{I} contributions to the correlation function are additive [10], that is,

$$G^{(1)}(\mathbf{x}, \mathbf{x}') = G_{\mathbf{C}}^{(1)}(\mathbf{x}, \mathbf{x}') + G_{\mathbf{I}}^{(1)}(\mathbf{x}, \mathbf{x}'), \quad (3.28)$$

where $G_{\mathbf{C}}^{(1)}$ and $G_{\mathbf{I}}^{(1)}$ are defined in Eqs. (3.19) and (B.16), respectively. It is interesting to note that $G_{\mathbf{C}}^{(1)}$ and $G_{\mathbf{I}}^{(1)}$ individually display an oscillatory decay behaviour — originating from the cutoff — an effect which correctly cancels when the two are added together.

Having calculated $g^{(1)}$, we obtain the coefficient α by fitting the algebraic decay law, Eq. (3.3), using nonlinear least squares; sample fits are shown in Fig. 3.4. The fit is conducted over the region between 10 and 40 de Broglie wavelengths. The short length scale cutoff is to avoid the contribution of the non-universal normal atoms, for which the thermal de Broglie

wavelength sets the appropriate decay length. The long distance cutoff is chosen to be small compared to the length scale L , to avoid the effect of periodic boundary conditions on the long range correlations.

The quality of the fitting procedure, and the breakdown of the expression in Eq. (3.3) at the BKT transition can be observed by adding an additional degree of freedom to the fitting function. In particular, at each temperature we fit the quadratic $\ln(g^{(1)}) = A - \tilde{\alpha} \ln(x) + \delta \ln^2(x)$ and extract the parameter δ ($\tilde{\alpha} \approx \alpha$ is discarded). The abrupt failure of the fits can be observed in the inset of Fig. 3.5 as a sudden increase in the value of $|\delta(T)|$ — an effect which is in excellent agreement with the value of T_{KT} as estimated from the superfluid fraction.




Figure 3.4: Sample fits to the algebraic decay of $g^{(1)}$ at multiples $T \approx 0.77, 0.93, 1.01$ and 1.12 of the transition temperature. High temperatures correspond to curves at the bottom of the figure which have rapid falloff of $g^{(1)}$ with distance. Fits are shown on a log-log scale in the inset to emphasise the failure of a power law in describing the behaviour of $g^{(1)}$ at high temperature.

Calculation of $g^{(1)}$ via interference patterns

So far a direct probe of the *in situ* spatial correlations has not been possible, although important progress has been made by the NIST group [23]. In the experiments of Hadzibabic *et al.* [60] a scheme proposed by Polkovnikov *et al.* [102] was used to infer these correlations from the “waviness” of interference patterns produced by pair of quasi-2D systems (see section 3.4.1).




Figure 3.5: Comparison of two methods for determining the algebraic decay coefficient $\alpha(T)$ for the first-order correlation function $g^{(1)}(\mathbf{x}, \mathbf{x}')$. The line with circle markers represents direct fits to $g^{(1)}$. These fits fail at the transition temperature as shown by the sharply diverging value of $|\delta(T)|$ in the inset. The filled points represent the values $\alpha'(T)$ obtained from a simulation of the experimental analysis procedure of Ref. [60], described in section 3.4.3. Horizontal dotted lines at 0.25 and 0.5 correspond to the expected values of α' just below and above the transition, respectively [60]. The vertical line is the BKT transition temperature, as estimated from the superfluid fraction calculated in section 3.4.2.

In this section we simulate the experimental data analysis method, and compare inferred predictions for the correlation function against those we can directly calculate. This allows us to characterise the errors associated with this technique arising from finite-size effects and finite expansion time.

To make this analysis we follow the procedure outlined in Ref. [60]. We fit our numerically generated interference patterns (see section 3.4.1) to the function

$$F(x, z) = G(z) \left[1 + c(x) \cos \left(\frac{2\pi z}{D} + \theta(x) \right) \right], \quad (3.29)$$

where $G(z)$ is a Gaussian envelope in the z -direction, $c(x)$ is the interference fringe contrast, D is the fringe spacing and $\theta(x)$ is the phase of the interference pattern in the z -direction.

Defining the function

$$C(L_x) = \frac{1}{L_x} \int_{-L_x/2}^{L_x/2} dx c(x) e^{i\theta(x)}, \quad (3.30)$$

the nature of spatial correlations is then revealed by the manner in which $\langle |C(L_x)|^2 \rangle$ decays with L_x . In particular, we identify the parameter α' , defined by $\langle |C(L_x)|^2 \rangle \propto L_x^{-2\alpha'}$ [102]. For an infinite 2D system in the superfluid regime ($T < T_{\text{KT}}^\infty$) $\alpha' = \alpha$ (i.e., α' corresponds to the algebraic decay of correlations). For $T > T_{\text{KT}}^\infty$, where correlations decay exponentially, α' is equal to 0.5.

Fitting $\langle |C(L_x)|^2 \rangle$ to the algebraic decay law $AL_x^{-2\alpha'}$ we can determine α' . A comparison between α' inferred from the interference pattern and α obtained directly from $g^{(1)}$ is shown in Fig. 3.5. Both methods give broadly consistent predictions for α when $T < T_{\text{KT}}$, however our results show that there is a clear quantitative difference between the two schemes, and that α' underestimates the coefficient of algebraic decay in the system (i.e., using α' in Eq. (3.4) would overestimate the superfluid density). Near and above the transition temperature, where the fits to $g^{(1)}$ fail, we observe that α' converges toward 0.5. The agreement between α and α' in the low temperature region improves as the size of the grid is increased.

3.4.4 Vortices and pairing

The simplest description of the BKT transition is that it occurs as a result of vortex pair unbinding: At $T < T_{\text{KT}}$ vortices only exist in pairs of opposite circulation, which unbind at the transition point to produce free vortices that destroy the superfluidity of the system. However, to date there are no direct experimental observations of this scenario, and theoretical studies of 2D Bose gases have been limited to qualitative inspection of the vortex distributions. In the c-field approach vortices and their dynamics are clearly revealed, unlike other ensemble-based simulation techniques such as quantum Monte Carlo where the vortices are obscured by averaging³. This gives us a unique opportunity to investigate the role of vortices and pairing in a dilute Bose gas.

We detect vortices in the c-field microstates by analysing the phase profile of the instantaneous field (see appendix B.3). An example of a phase profile of a field for $T < T_{\text{KT}}$ is shown in Fig. 3.6(a). The vortex locations reveal a pairing character, that is, the close proximity of pairs of positive (clockwise) and negative (counterclockwise) vortices relative to the average vortex separation. An important qualitative feature of our observed vortex distributions is that at high temperatures, pairing does not disappear from the system entirely. Indeed,

³For example, [42] calculates the vortex density, but only indirectly via a relation with the quasiparticle density.

most vortices at high temperature could be considered paired or grouped in some manner, as shown in Fig. 3.6(b). Perhaps this is not surprising, since positive and negative vortices have a logarithmic attraction, and we observe them to create and annihilate readily in the c-field dynamics. However, this does indicate that the use of pairing to locate the transition may be ambiguous, and we examine this aspect further below.

It is also of interest to measure the number of vortices, N_v , present in the system as a function of temperature (see Fig. 3.7). At the lowest temperatures the system is in an ordered state, and the energetic cost of having a vortex is prohibitive. As the temperature increases there is a rapid growth of vortex population leading up to the transition point followed by linear growth above T_{KT} .

Radial vortex density

The most obvious way to characterise vortex pairing is by defining a pair distribution function for vortices of opposite sign. Adopting the notation of Ref. [51], this is

$$G_{v,\pm}^{(2)}(\mathbf{r}) = \langle \rho_{v,+}(\mathbf{0}) \rho_{v,-}(\mathbf{r}) \rangle, \quad (3.31)$$

where $\rho_{v,+}$ is the vortex density function which consists of a sum of delta spikes,

$$\rho_{v,+}(\mathbf{r}) = \sum_{i=1}^{N_{v,+}} \delta(\mathbf{r} - \mathbf{r}_i^+) \quad (3.32)$$

for positive vortices at positions $\{\mathbf{r}_i^+\}$. We use the analogous definition for $\rho_{v,-}$. The associated dimensionless two-vortex correlation function is

$$g_{v,\pm}^{(2)}(\mathbf{r}) = \frac{G_{v,\pm}^{(2)}(\mathbf{r})}{\langle \rho_{v,+}(\mathbf{0}) \rangle \langle \rho_{v,-}(\mathbf{r}) \rangle}. \quad (3.33)$$

The angular average of $g_{v,\pm}^{(2)}$ can be calculated directly from the detected vortex positions using a binning procedure on the pairwise distances $\|\mathbf{r}_i^+ - \mathbf{r}_j^-\|$, and is shown in Fig. 3.8.

These results quantify the effect discussed earlier: Positive and negative vortices show a pairing correlation that does not disappear above T_{KT} . The characteristic size of this correlation, given by twice the width of the peak feature in Fig. 3.8, is $l_{\text{cor}} \sim 3\mu\text{m}$ (taking full width half maximum).

The shape of our pairing peak is qualitatively similar to that described in Ref. [51]. However, in contrast to their results the width does not appear to change appreciably with temperature. Additional simulations show that increasing the interaction strength causes the peak to become

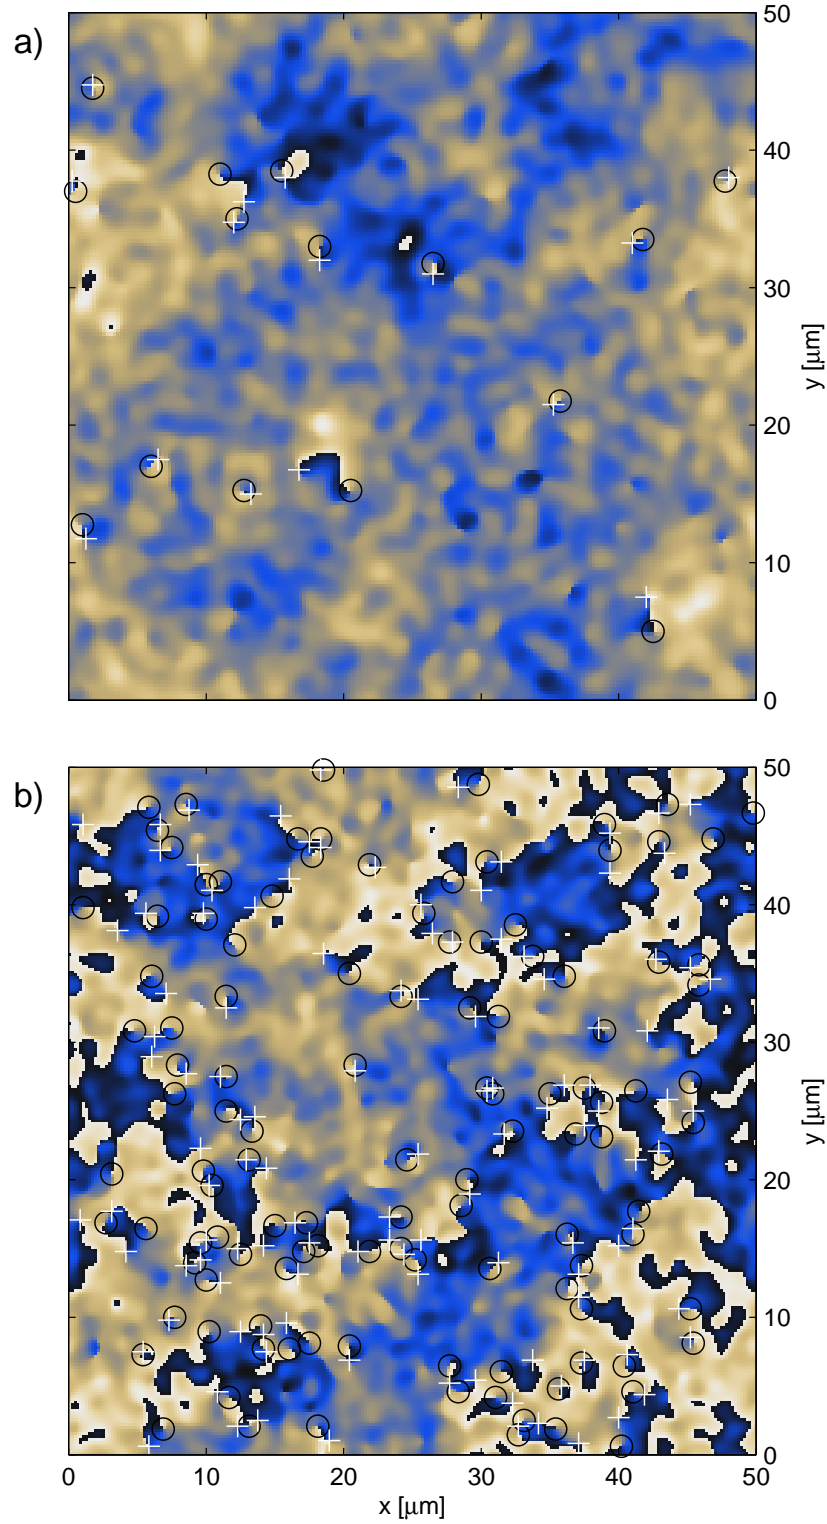


Figure 3.6: Phase profile of a c-field with vortices indicated. Vortices with clockwise (white +) and anticlockwise (black o) circulation. The phase of the classical field is indicated by shading the background between dark blue (phase 0) and light yellow (phase 2π). (a) Distinctive pairing below the transition at $T = 207\text{nK} \approx 0.93T_{\text{KT}}$. (b) A “vortex plasma” above the transition at $T = 238\text{nK} \approx 1.07T_{\text{KT}}$.




Figure 3.7: Total number of vortices (dots) and number of unpaired vortices (circles) as a function of temperature near the transition. While N_v at the transition temperature is already very high, N_u becomes nonzero only close to the transition, providing clear evidence of vortex unbinding at work. The inset shows the variation in the total number over the full temperature range of the simulations. Above the transition temperature the growth in the number of vortices becomes linear with temperature.

squarer and wider. It is clear that while the pair size and strength revealed in $g_{v,\pm}^{(2)}(r)$ does not change appreciably as the transition is crossed, the amount of pairing relative to the background uncorrelated vortices changes considerably. This background of uncorrelated vortices is given by the horizontal plateau $g_{v,\pm}^{(2)}(r) \rightarrow 1$ at large r as shown in the inset.

Revealing unpaired vortices with coarse-graining

The function $G_{v,\pm}^{(2)}(r)$ clearly indicates the existence of vortex pairing in the system. However, it does not provide a convenient way to locate the transition temperature, because a large amount of pairing exists both below and above the transition: The expected number of neighbours for any given vortex — roughly, the area of the pairing peak of $\langle n_{v,+} \rangle G_{v,\pm}^{(2)}(r)$ shown in Fig. 3.8 — does not change dramatically across the transition. $\langle n_{v,+} \rangle = \langle n_v \rangle / 2$ is the expected density of positive vortices.

We desire a quantitative observation of vortex unbinding at the transition and have therefore investigated several measures of vortex pairing⁴. However, measures based directly on the full

⁴For example, the Hausdorff distance (see, for example, [94, p. 105]) between the set $\{\mathbf{r}_i^+\}$ of positive vortices




Figure 3.8: Angular average of the two-vortex pair distribution functions for vortices of opposite sign. Three temperatures centred about the transition are shown: dot markers $T = 194\text{nK} \approx 0.9T_{\text{KT}}$, $f_c = 0.34$; circle markers $T = 217\text{nK} \approx 1.01T_{\text{KT}}$, $f_c = 0.076$; cross markers $T = 236\text{nK} \approx 1.1T_{\text{KT}}$, $f_c = 0.006$. The vertical dotted line shows the value of the healing length at $T = 0$. The main plot shows $g_{v,\pm}^{(2)}$ normalised by the positive vortex density; comparable magnitudes for the peaks near $r = 0$ show that vortex pairing remains important over the range of temperatures studied, not only below the transition. The inset shows $g_{v,\pm}^{(2)}$ in the natural dimensionless units for which $g_{v,\pm}^{(2)}(r) \rightarrow 1$ as $r \rightarrow \infty$.

set of vortex positions seem to suffer from the proliferation of vortices at high temperature — an effect that tends to wash out clear signs of vortex unbinding. With this in mind, we have developed a procedure for measuring the number of *unpaired* vortices in our simulations, starting from the classical field rather than the full set of vortex positions.

The basis of our approach for detecting unpairing is to coarse-grain the classical field by convolution with a Gaussian filter of spatial width (standard deviation) σ_f . This removes all vortex pairs on length scales smaller than σ_f . Figure 3.9 shows the count of remaining vortices as a function of filter width, along with some examples of coarse-grained fields. For $\sigma_f \gtrsim l_{\text{cor}}$, the number of remaining vortices levels off and only decreases slowly with increasing σ_f . Ultimately the number of remaining vortices goes to zero as $\sigma_f \rightarrow L$.

Setting the filter width to be larger than the characteristic pairing distance, l_{cor} , yields a

and the set $\{\mathbf{r}_i^-\}$ of negative vortices.

coarse-grained field from which the pairs have been removed, but unpaired vortices remain. In our simulations we have $l_{\text{cor}} \approx 3 \mu\text{m}$; we take the vortices that remain after coarse-graining with a Gaussian of standard deviation $\sigma_f = 5 \mu\text{m}$ to give an estimate of the number of unpaired vortices, N_u . Figure 3.7 shows that N_u becomes nonzero only near the transition, in contrast to N_v which is nonzero well below T_{KT} . The sharp increase in N_u at T_{KT} is a quantitative demonstration of vortex unbinding at work.




Figure 3.9: The coarse-graining procedure: number of vortices as a function of filter width for a temperature near the transition. The smooth curve is an average over many realisations of the field, whereas the stepped curve shows typical behaviour of the number for a single field. Insets show the coarse-grained fields for various filter widths; the transformation removes vortex-antivortex pairs that are separated by approximately less than the standard deviation of the filter. In this example $N_u = 4$ unpaired vortices remain at $\sigma_f = 5 \mu\text{m}$.

In the experiment of Ref. [60], the fraction of interference patterns with dislocations (see, for example, Figs. 3.1(c) and (d)) was measured. While isolated vortices are clearly identified by interference pattern dislocations, a lack of spatial resolution in experiments means that this type of detection method obscures the observation of tightly bound vortex pairs. The experimental resolution of $3 \mu\text{m}$ is broadly consistent with the scale of the coarse-graining filter (i.e., $\sigma_f = 5 \mu\text{m}$). With this in mind, we introduce the quantity $p_u(T)$, defined as the probability of observing an unpaired vortex in a $50 \times 50 \mu\text{m}$ control volume at a given temperature⁵. For

⁵We choose a fixed control volume with $L = 50 \mu\text{m}$ in order to compare results between simulations with

the 50 μm grid we have simply $p_u(T) = \text{Pr}(N_u \geq 1)$.

Computing $p_u(T)$ from our simulations yields the results shown in Fig. 3.10. Our results show a dramatic jump in p_u at a temperature that is consistent with the transition temperature T_{KT} determined from the superfluid fraction calculation presented in section 3.4.2.




Figure 3.10: Comparison of vortex unpairing measures. The dots are our pairing measure based on coarse-graining the field. Circles represent the pairing as determined by the number of dislocations in the simulated interference patterns. This was the same method used in the experimental analysis of Ref. [60] and coincides remarkably well with our coarse-graining based measure. Both curves are consistent with the vertical line showing the transition temperature T_{KT} as determined from the superfluid fraction calculation in section 3.4.2. The inset shows the calculated coarse-grained pairing measure for all three grid sizes, along with vertical lines showing the estimates for T_{KT} derived from the superfluid fraction calculations.

From the definition, we expect that p_u should be close to the experimentally measured frequency of dislocations. To demonstrate this relationship, we have simulated interference patterns (as described in section 3.4.1) and detected dislocations using the experimental procedure of Ref. [60]: A phase gradient $d\theta/dx$ was considered to mark a dislocation whenever $|d\theta/dx| > \pi/4$ rad/ μm . From this we can compute the probability of detecting at least one dislocation as a function of temperature. As shown in Fig. 3.10, the results of this procedure compare very favourably with our measure of pairing based on p_u . We note that inhomoge-

different grid sizes.

neous effects in experiments probably broaden the jump in p_u appreciably compared to our homogeneous results.

3.5 Conclusion

In this chapter we have used c-field simulations of a finite-sized homogeneous system in order to investigate the physics of the 2D Bose gas in a regime corresponding to current experiments. We have directly computed the condensate and superfluid fractions as a function of temperature, and made comparisons to the superfluid fraction inferred both from the first-order correlation function, and the experimental interference scheme. Our results for these quantities provide a quantitative test of the interference scheme for a finite system.

An intriguing possibility is the direct experimental observation of vortex-antivortex pairs, their distribution in the system, and hence a quantitative measurement of their unbinding at the BKT transition. We have calculated the vortex correlation function across the transition and provided a coarse-graining scheme for distinguishing unpaired vortices. These results suggest that the dislocations observed in experiments, due to limited optical resolution, provide an accurate measure of the unpaired vortex population and accordingly are a strong indicator of the BKT transition.

We briefly discuss the effect that harmonic confinement (present in experiments) would have on our predictions. The spatial inhomogeneity will cause the superfluid transition to be gradual, occurring first at the trap centre where the density is highest, in contrast to our results where the transition occurs in the bulk.

Bisset *et al.* [13] used an extension of the c-field method for the trapped 2D gas to examine $g^{(1)}$ and found similar results for the onset of algebraic decay of correlations at the transition. Their analysis was restricted to the small region near the trap centre where the density is approximately constant; we expect the results of our vortex correlation function and the coarse-graining scheme should similarly be applicable to the trapped system in the central region. Except in very weak traps, the size of this region is relatively small and will likely prove challenging to measure experimentally.

Our results for the homogeneous gas emphasise the clarity with which *ab initio* theoretical methods can calculate quantities directly observable in experiments, such as interference patterns. This should allow direct comparisons with experiments, providing stringent tests of many-body theory.

Superfluid fraction and the PGPE

In this chapter we describe a method for calculating the superfluid fraction from a PGPE classical field simulation. We first present the analytical derivation, showing how the superfluid density arises in certain limits of the momentum density autocorrelations. This is followed by discussion of a numerical implementation relevant to the 2D simulations of chapter 3.

4.1 Introduction

Superfluidity is a famous example of macroscopic quantum behaviour, and is typically discussed in macroscopic terms. In particular, one characterises a superfluid by its zero viscosity; the ability to “flow without friction” through a narrow channel. Although such macroscopic ideas are easily expressed, it is not trivial to connect them to microscopic theories such as the PGPE formalism in an efficient way. In the following we briefly provide some relevant background before following with the details of our derivation in the next section. For further background theory we refer the reader to chapter 6 of Ref. [100] which provides an accessible overview of superfluid theory as relevant to experiments on ultracold Bosons.

According to Landau’s phenomenological model of superfluidity [78], it is possible to model a superfluid system as a “mixture” of two liquids: a superfluid part without viscosity, and a normal part. This idea was introduced to describe the residual viscosity which remains in liquid helium, even below the superfluid transition temperature. One defines a *superfluid fraction* as the ratio $f_s \equiv \rho_s/\rho$ of the superfluid density ρ_s to the total density ρ of the system¹.

The two-fluid model was put on firmer ground by Putterman and Roberts in Ref. [108]. Starting from the equations for a single nonlinear classical fluid, they considered the presence of small amplitude excitations on top of a background fluid. Using only a separation of scales argument, they have derived kinetic equations for these thermal excitations. In the hydrodynamic (collision dominated) regime, the model then reduces to the Landau two-fluid model

¹Note that we follow convention and use the *mass* density ρ in the current chapter rather than the number density n that is used in the rest of the thesis.

of superfluidity. The connection between this model and the classical field methods is further described in [112].

In Landau's theory, superfluidity may be predicted from the form of the energy spectrum of elementary excitations. Let $\epsilon(\mathbf{p})$ be the energy of an excitation with momentum \mathbf{p} measured with respect to a stationary background of fluid. To understand the origin of superfluidity, we imagine that the fluid occupies a narrow channel with walls moving at velocity \mathbf{u} with respect to the fluid. In the frame of reference where the walls are stationary, the energy of the excitation is found to be $\epsilon'(\mathbf{p}) \equiv \epsilon(\mathbf{p}) - \mathbf{u} \cdot \mathbf{p}$ after applying a Galilean transformation to the Hamiltonian². This is negative for sufficiently large \mathbf{u} , making the formation of excitations energetically favourable. However, if \mathbf{u} is small enough, it may be the case that $\epsilon'(\mathbf{p}) > 0$ for all \mathbf{p} and excitations are energetically forbidden. This leads to Landau's criterion for superfluidity: $\|\mathbf{u}\|$ must be smaller than the critical velocity,

$$v_c \equiv \min_{\mathbf{p}} \frac{\epsilon(\mathbf{p})}{p}, \quad (4.1)$$

where $p = \|\mathbf{p}\|$.

It is worth noting that any system with quadratic dispersion relation $\epsilon(\mathbf{p}) \propto p^2$ for small p has a critical velocity of zero and cannot be a superfluid. In particular, this includes the ideal gas, where the elementary excitations are simply particles and we have $\epsilon(\mathbf{p}) = p^2/2m$. On the other hand, introducing interactions as in the Hamiltonian Eq. (2.47) modifies the energy so that $\epsilon(\mathbf{p}) \propto p$ at small p , which allows the system to support superfluidity. For sufficiently weak interactions, Hamiltonian (2.47) may be approximately diagonalised via the Bogoliubov transformation (see, for example, [100, §4.3]), yielding the classic dispersion relation

$$\epsilon(\mathbf{p}) = \sqrt{\frac{U_0 n}{m} p^2 + \left(\frac{p^2}{2m}\right)^2} \quad (4.2)$$

for the energies of elementary excitations, known as Bogoliubov quasiparticles.

The approximate diagonalisation discards terms corresponding to the interaction of quasiparticles, which is a good approximation for sufficiently weak interactions and low temperatures. Further, this leads to a well known method for computing the superfluid fraction. At nonzero temperature, the Bose distribution gives the number of non-interacting quasiparticles at each energy,

$$N_{\mathbf{p}} = \left[\exp \left(\frac{\epsilon(\mathbf{p}) + \mathbf{p} \cdot \mathbf{u}}{kT} \right) - 1 \right]^{-1}. \quad (4.3)$$

²Note that in our notation \mathbf{u} is the velocity of the walls with respect to the fluid, so the superfluid velocity is $-\mathbf{u}$ with respect to the walls. This is opposite from the convention used in Ref. [100].

This may be used to evaluate the expected momentum density $\langle \hat{\mathbf{p}} \rangle_{\mathbf{u}}$, which we attribute to the motion of the normal fraction, so that

$$\rho_n \mathbf{u} = \langle \hat{\mathbf{p}} \rangle_{\mathbf{u}} = \int \mathbf{p} N_{\mathbf{p}} \frac{d\mathbf{p}}{h^3}, \quad (4.4)$$

where we have used the phase space volume $1/h^3$ appropriate to three dimensions. For consistency with the next section, we use the notation $\langle \cdot \rangle_{\mathbf{u}}$ to mean an expectation value with respect to the system where the walls are moving. After some manipulation, we arrive at an expression for the density of the normal fraction [100]

$$\rho_n = -\frac{1}{3} \int \frac{dN_{\mathbf{p}}(\epsilon)}{d\epsilon} p^2 \frac{d\mathbf{p}}{h^3}. \quad (4.5)$$

This method is valid when the physical picture of non-interacting quasiparticles is valid — in particular, one requires weak interactions and low temperatures. In the context of the classical field method for ultracold Bose gases, the method has previously been used to compute the superfluid fraction, see for example Ref. [127]. Unfortunately it is not valid for the system considered in chapter 3 for two reasons. First, we wish to compute the superfluid fraction over a wide range of temperatures, from zero to slightly above the transition and the underlying assumptions are invalid near the transition [100, p. 66]. Second, the weakly interacting limit is especially difficult to reach in two dimensions because it requires the inequality $\ln \ln(1/na^2) \ll 1$ to be satisfied as discussed in Ref. [43] (see also Ref. [106]).

Given our two-dimensional system at moderate temperatures, we require a non-perturbative alternative to Eq (4.5). Starting from the macroscopic definition of superfluidity, one naturally imagines performing a time-varying numerical experiment in order to determine the superfluid density within the PGPE formalism. For example, in the homogeneous case we might examine the drag force produced when a perturbing potential is moved across the system; zero drag would imply a superfluid fraction of 100%. Dynamical PGPE simulations were used in a similar way in Ref. [119] to provide evidence for superfluidity in a 2D trapped Bose gas, by analysing the dynamics of the “scissors mode” oscillation.

Deducing superfluidity from dynamical simulations is certainly possible, but is far from ideal: For the work presented in chapter 3 it would mean performing an additional set of numerically expensive simulations at every temperature of interest. Our method address this problem by using linear response theory to relate the superfluid fraction to the long wavelength limit of the second order momentum density correlations. The method is attractive because the momentum correlations may be extracted directly from PGPE simulations at thermal equilibrium. This allows the superfluid fraction to be computed from the same set of simulations

as the temperature, chemical potential, and other thermal properties; there is no need to perform an expensive special purpose simulation for the sole purpose of calculating the superfluid fraction.

4.2 Momentum density correlations and the superfluid fraction

Our derivation is based on the microscopic theory presented in Ref. [45, p.214], (see also [8] and [100, p.96]). The central idea is to establish a relationship between (i) the autocorrelations of the momentum density in the simulated ensemble and (ii) the linear response of the fluid to slowly moving solid boundaries; (i) is a quantity we can calculate, while (ii) is related to the basic properties of a superfluid via a simple thought experiment.

To connect the macroscopic, phenomenological description of superfluidity with our microscopic theory, we make use of the standard thought experiment shown schematically in Fig. 4.1(b): Consider an infinitely long box, B containing superfluid, and accelerate the box along its long axis until it reaches a small velocity \mathbf{u} . Due to viscous interactions with the walls, such a box filled with a normal fluid should have a momentum density at equilibrium of $\langle \hat{\mathbf{p}} \rangle_{\mathbf{u}} = \rho \mathbf{u}$. As above, the notation $\langle \cdot \rangle_{\mathbf{u}}$ denotes an expectation value in the ensemble with walls moving with velocity \mathbf{u} .

Because the superfluid part is nonviscous, the observed value for the momentum density in a superfluid is less than the value $\rho \mathbf{u}$ expected for a classical fluid. As above, we attribute the observed momentum density, $\rho_n \mathbf{u}$, to the “normal fraction” where ρ_n is the normal fluid density. The superfluid fraction remains stationary in the lab frame, even at equilibrium and makes up the remaining mass with density $\rho_s = \rho - \rho_n$.

In order to apply the usual procedures of statistical mechanics to the thought experiment, we consider two frames: the “lab frame” in which the walls move with velocity \mathbf{u} in the x -direction and the “wall frame” in which the walls are at rest.

Assuming that the fluid is in thermal equilibrium with the walls, the density matrix in the grand canonical ensemble is given by the usual expression $\hat{\rho} = e^{-\beta(\hat{H}_{\mathbf{u}} - \mu \hat{N})} / \text{Tr} (e^{-\beta(\hat{H}_{\mathbf{u}} - \mu \hat{N})})$ where $\hat{H}_{\mathbf{u}}$ is the Hamiltonian of the system in the wall frame and $\beta = 1/k_B T$. A Galilean transformation relates $\hat{H}_{\mathbf{u}}$ to the Hamiltonian in the lab frame, $\hat{H}_{\mathbf{u}} = \hat{H} - \mathbf{u} \cdot \hat{\mathbf{P}} + \frac{1}{2} M u^2$, where $\hat{\mathbf{P}} = \int_B d^2 \mathbf{x} \hat{\mathbf{p}}(\mathbf{x})$ is the total momentum, $M = mN$ is the total mass and $\hat{\mathbf{p}}(\mathbf{x})$ is the momentum density operator at point \mathbf{x} . The expectation value for the momentum density in

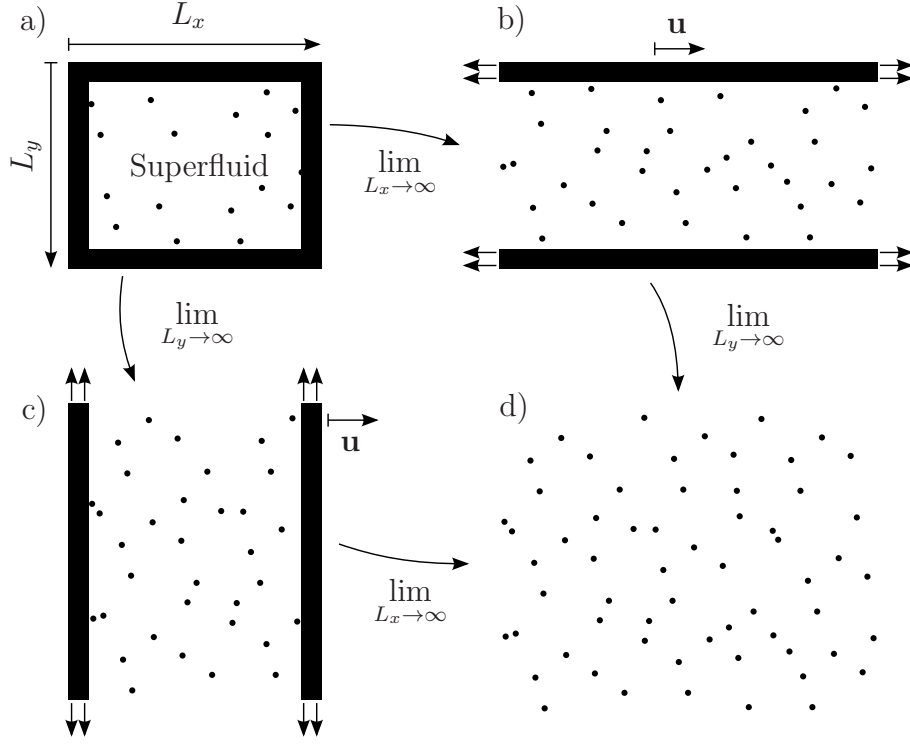


Figure 4.1: Thought experiment used in deriving the superfluid density. The walls move with velocity \mathbf{u} in the x direction. To begin with, we imagine that the superfluid sits in a box of dimensions $L_x \times L_y$ as shown in (a). We later take the limit as the box walls recede to infinity to get the thermodynamic limit (d). The order of the limits is critically important: the path (b) leads to superflow while the path (c) results in the entire fluid moving along with the walls.

the presence of moving walls is then given by the expression

$$\langle \hat{\mathbf{p}}(\mathbf{x}) \rangle_{\mathbf{u}} = \text{Tr}[\hat{\rho} \hat{\mathbf{p}}(\mathbf{x})], \quad (4.6)$$

$$= \frac{\text{Tr} \left(e^{-\beta(\hat{H} - \mathbf{u} \cdot \hat{\mathbf{P}} + (mu^2/2 - \mu)\hat{N})} \hat{\mathbf{p}}(\mathbf{x}) \right)}{\text{Tr} \left(e^{-\beta(\hat{H} - \mathbf{u} \cdot \hat{\mathbf{P}} + (mu^2/2 - \mu)\hat{N})} \right)}. \quad (4.7)$$

Expanding this expression to first order in \mathbf{u} yields

$$\langle \hat{\mathbf{p}}(\mathbf{x}) \rangle_{\mathbf{u}} = \langle \hat{\mathbf{p}}(\mathbf{x}) \rangle + \beta \left(\langle \hat{\mathbf{p}}(\mathbf{x}) \mathbf{u} \cdot \hat{\mathbf{P}} \rangle - \langle \hat{\mathbf{p}}(\mathbf{x}) \rangle \langle \mathbf{u} \cdot \hat{\mathbf{P}} \rangle \right), \quad (4.8)$$

where all the expectation values on the right hand side are now taken in the *equilibrium* ensemble with the walls at rest. Since $\langle \hat{\mathbf{p}}(\mathbf{x}) \rangle = 0$ in our equilibrium ensemble, this simplifies to

$$\langle \hat{\mathbf{p}}(\mathbf{x}) \rangle_{\mathbf{u}} = \beta \langle \hat{\mathbf{p}}(\mathbf{x}) \hat{\mathbf{P}} \rangle \cdot \mathbf{u}, \quad (4.9)$$

$$= \beta \int_B d^2 \mathbf{x}' \langle \hat{\mathbf{p}}(\mathbf{x}) \hat{\mathbf{p}}(\mathbf{x}') \rangle \cdot \mathbf{u}, \quad (4.10)$$

where $\hat{\mathbf{p}}(\mathbf{x})\hat{\mathbf{p}}(\mathbf{x}')$ is a rank-two tensor; the outer product of $\hat{\mathbf{p}}(\mathbf{x})$ and $\hat{\mathbf{p}}(\mathbf{x}')$. [In two dimensions this means $\langle\hat{\mathbf{p}}(\mathbf{x})\hat{\mathbf{p}}(\mathbf{x}')\rangle$ is a 2×2 matrix for each pair of coordinates $(\mathbf{x}, \mathbf{x}')$.]

To make further progress, we wish to take the limit as the system gets very large (we will notate this limit as $B \rightarrow \infty$). To this end, we first consider some properties of the correlation functions in the infinite system. The infinite system is homogeneous, which implies that $\langle\hat{\mathbf{p}}(\mathbf{x})\hat{\mathbf{p}}(\mathbf{x}')\rangle_\infty = \langle\hat{\mathbf{p}}(\mathbf{x}+\mathbf{r})\hat{\mathbf{p}}(\mathbf{x}'+\mathbf{r})\rangle_\infty$ for any \mathbf{r} , where $\langle\cdot\rangle_\infty$ indicates an average in the infinite system. As a consequence, we may express the correlations — in the infinite system — in terms of the Fourier transform in the relative coordinate $\mathbf{x}' - \mathbf{x}$:

$$\langle\hat{\mathbf{p}}(\mathbf{x})\hat{\mathbf{p}}(\mathbf{x}')\rangle_\infty = \langle\hat{\mathbf{p}}(\mathbf{0})\hat{\mathbf{p}}(\mathbf{x}' - \mathbf{x})\rangle_\infty \quad (4.11)$$

$$= \frac{1}{(2\pi)^2} \int d^2\mathbf{k} e^{i\mathbf{k}\cdot(\mathbf{x}'-\mathbf{x})} \chi(\mathbf{k}), \quad (4.12)$$

where all the important features of the correlations are now captured by the tensor

$$\chi(\mathbf{k}) = \int d^2\mathbf{r} e^{-i\mathbf{k}\cdot\mathbf{r}} \langle\hat{\mathbf{p}}(\mathbf{0})\hat{\mathbf{p}}(\mathbf{r})\rangle_\infty. \quad (4.13)$$

Because of the isotropy of the fluid in the infinite system, $\chi(\mathbf{k})$ obeys the transformation law $\chi(O\mathbf{k}) = O^{-1}\chi(\mathbf{k})O$, for any 2×2 orthogonal matrix O . This implies that χ may be decomposed into the sum of longitudinal and transverse parts:

$$\chi(\mathbf{k}) = \tilde{\mathbf{k}}\tilde{\mathbf{k}}\chi_l(k) + (I - \tilde{\mathbf{k}}\tilde{\mathbf{k}})\chi_t(k) \quad (4.14)$$

where $\tilde{\mathbf{k}} = \mathbf{k}/k$, $k = \|\mathbf{k}\|$, I is the identity and the juxtaposition of vectors $\tilde{\mathbf{k}}\tilde{\mathbf{k}}$ represents the outer product as above. The transverse and longitudinal functions χ_t and χ_l are scalars that depend *only* on the length k .

We now return our attention to the finite system. If the finite box B is large then the momentum correlations in the bulk will be very similar to the values for the infinite system. Therefore, when \mathbf{x} and \mathbf{x}' are far from the boundaries, we may approximate

$$\langle\hat{\mathbf{p}}(\mathbf{x})\hat{\mathbf{p}}(\mathbf{x}')\rangle \approx \langle\hat{\mathbf{p}}(\mathbf{x})\hat{\mathbf{p}}(\mathbf{x}')\rangle_\infty \quad (4.15)$$

$$= \frac{1}{(2\pi)^2} \int d^2\mathbf{k} e^{i\mathbf{k}\cdot(\mathbf{x}'-\mathbf{x})} \chi(\mathbf{k}) \quad (4.16)$$

which in combination with Eq. (4.10) yields

$$\langle \hat{\mathbf{p}}(\mathbf{x}) \rangle_{\mathbf{u}} \approx \beta \int_B d^2 \mathbf{x}' \frac{1}{(2\pi)^2} \int d^2 \mathbf{k} e^{i\mathbf{k} \cdot (\mathbf{x}' - \mathbf{x})} \chi(\mathbf{k}) \cdot \mathbf{u} \quad (4.17)$$

$$= \beta \int d^2 \mathbf{k} \Delta_B(\mathbf{k}) e^{-i\mathbf{k} \cdot \mathbf{x}} \chi(\mathbf{k}) \cdot \mathbf{u}. \quad (4.18)$$

Here we have defined the nascent delta function $\Delta_B(\mathbf{k}) \equiv \frac{1}{(2\pi)^2} \int_B d^2 \mathbf{x}' e^{i\mathbf{k} \cdot \mathbf{x}'}$ which has the property $\Delta_B(\mathbf{k}) \rightarrow \delta(\mathbf{k})$ as $B \rightarrow \infty$.

We are now in a position to carry out the limiting procedure to increase the box size to infinity. However, care must be taken because the simple expression $\lim_{B \rightarrow \infty} \langle \hat{\mathbf{p}}(\mathbf{x}) \rangle_{\mathbf{u}}$ is not well defined without further qualification of the limiting process $B \rightarrow \infty$.

To resolve this subtlety we must insert a final vital piece of physical reasoning. Let us assume for simplicity that \mathbf{u} is directed along the x -direction, and the box B is aligned with the x and y axes with dimensions $L_x \times L_y$. As shown in Fig. 4.1, there are two possibilities for taking the limits, representing different physical situations.

On the one hand [Fig. 4.1(b)], we may take the limit $L_x \rightarrow \infty$ first, which gives us an infinitely long channel in which superfluid can remain stationary while only the normal fraction moves with the walls in the x -direction. We have

$$\rho_n \mathbf{u} = \lim_{L_y \rightarrow \infty} \lim_{L_x \rightarrow \infty} \langle \hat{\mathbf{p}}(\mathbf{x}) \rangle_{\mathbf{u}} \quad (4.19)$$

$$= \lim_{L_y \rightarrow \infty} \lim_{L_x \rightarrow \infty} \beta \int d^2 \mathbf{k} \Delta_B(\mathbf{k}) e^{-i\mathbf{k} \cdot \mathbf{x}} \chi(\mathbf{k}) \cdot \mathbf{u} \quad (4.20)$$

$$= \beta \lim_{k_y \rightarrow 0} \lim_{k_x \rightarrow 0} \chi(\mathbf{k}) \cdot \mathbf{u} \quad (4.21)$$

where we use the fact that $\Delta_B(\mathbf{k})$ can be decomposed into the product $\Delta_{L_x}(k_x) \Delta_{L_y}(k_y)$ with $\Delta_L(k) \rightarrow \delta(k)$ as $L \rightarrow \infty$. Employing the decomposition of χ given in Eq. (4.14) allows the density of the normal fraction to be related to the transverse component of χ evaluated at zero:

$$\rho_n = \beta \lim_{k \rightarrow 0} \chi_t(k) = \beta \chi_t(0). \quad (4.22)$$

On the other hand [Fig. 4.1(c)] we may take the limit $L_y \rightarrow \infty$ first, resulting in an infinitely long channel — with velocity *perpendicular* to the walls — in which the entire body of the fluid must move regardless of whether it is a superfluid or not. In a similar way to the previous paragraph, $\rho \mathbf{u} = \beta \lim_{k_x \rightarrow 0} \lim_{k_y \rightarrow 0} \chi(\mathbf{k}) \cdot \mathbf{u}$, and making use of the decomposition in Eq. (4.14) we find that the total density is related to the longitudinal component of the correlations:

$$\rho = \beta \lim_{k \rightarrow 0} \chi_l(k) = \beta \chi_l(0). \quad (4.23)$$

With these expressions, the normal fraction f_n may finally be expressed directly as

$$f_n = \rho_n / \rho = \lim_{k \rightarrow 0} \chi_t(k) / \lim_{k \rightarrow 0} \chi_l(k) \quad (4.24)$$

while the superfluid fraction is $f_s = 1 - f_n$. Thus, we have expressed the superfluid and normal fractions in terms of a correlation function which can be directly computed in the thermal ensemble; there is no need to deal with difficult moving boundary conditions or other dynamical perturbations in the simulation itself.

4.3 Numerical procedure

To determine the superfluid fraction from a PGPE simulation, we need to estimate the tensor of momentum density correlations χ using simulation results. For a finite system constrained to a periodic simulation box — as studied in chapter 3 — we may only compute the momentum correlations at discrete grid points. The discrete analogue of Eq. (4.13) leads to the expression

$$\chi(\mathbf{k}) \propto \langle \mathbf{p}_{\mathbf{k}} \mathbf{p}_{-\mathbf{k}} \rangle \quad (4.25)$$

where the constant of proportionality is not important to the final result, and $\mathbf{p}_{\mathbf{k}}$ are the discrete Fourier coefficients of $\mathbf{p}(\mathbf{x})$ over our simulation box.

The momentum density operator is given by

$$\hat{\mathbf{p}}(\mathbf{x}) = \frac{i\hbar}{2} \left[(\nabla \hat{\psi}^\dagger(\mathbf{x})) \hat{\psi}(\mathbf{x}) - \hat{\psi}^\dagger(\mathbf{x}) \nabla \hat{\psi}(\mathbf{x}) \right] \quad (4.26)$$

which may be derived by considering the continuity equation for the number density, $\langle \hat{\psi}^\dagger(\mathbf{x}) \hat{\psi}(\mathbf{x}) \rangle$. For a given classical field Eq. (3.9), the Fourier coefficients of \mathbf{p} may be written as

$$\mathbf{p}_{\mathbf{k}} = \frac{\hbar}{2\sqrt{A_B}} \sum_{\mathbf{k}'} (2\mathbf{k}' + \mathbf{k}) c_{\mathbf{k}'}^* c_{\mathbf{k}+\mathbf{k}'}, \quad (4.27)$$

where A_B is the area of the system. Computing a value for all $\mathbf{p}_{\mathbf{k}}$ at each time step, we then evaluate $\chi(\mathbf{k})$ via the usual ergodic averaging procedure using Eq. (4.25).

Having evaluated $\chi(\mathbf{k})$, we are left with performing the decomposition into longitudinal and transverse parts. For this, simply note that Eq. (4.14) implies $\chi_l(k) = \tilde{\mathbf{k}} \cdot \chi(\mathbf{k}) \cdot \tilde{\mathbf{k}}$, and $\chi_t(k) = \tilde{\mathbf{w}} \cdot \chi(\mathbf{k}) \cdot \tilde{\mathbf{w}}$, where $\tilde{\mathbf{w}}$ is a unit vector perpendicular to $\tilde{\mathbf{k}}$.

Values for χ_t and χ_l may be collected for all angles as a function of k , and a fitting procedure used to perform the extrapolation $k \rightarrow 0$; this procedure is illustrated in Fig. 4.2. At low

temperatures, the extrapolation is quite reliable, but becomes more difficult near the superfluid transition where sampling noise increases and $\chi_t(k)$ changes rapidly near $k = 0$. Without a known functional form, we settled for a quadratic weighted least squares fit of $\ln(\chi_t)$ and $\ln(\chi_l)$ versus k . A weighting of $1/k$ was used to account for the fact that the density of samples of χ vs k scales proportionally with k due to the square grid on which $\chi(\mathbf{k})$ is evaluated. The logarithm was used to improve the fits of χ_t very near the transition where it varies non-quadratically near $k = 0$. The fitting procedure and extrapolation to $k = 0$ generally produces reasonable results, but is somewhat sensitive to numerical noise. For this reason, the computed superfluid fraction at high temperatures is not exactly zero (see Fig. 3.2 on page 49).




Figure 4.2: Example fitting and extrapolation to $k = 0$ for the transverse and longitudinal components of the momentum density autocorrelation tensor, χ . The apparent functional form for χ_t and χ_l changes with temperature — particularly near the transition — which along with the sampling noise makes them difficult to fit reliably. The data shown corresponds to a temperature $T \approx 0.99T_{\text{KT}}$ slightly below the transition.

4.4 Discussion

We have applied the technique described in this chapter to study the BKT phase in chapter 3. This included computing the superfluid fraction as a function of temperature — see Fig. 3.2 on page 49 — and the shape of the curve is consistent with the expectation of a universal

jump³ in the superfluid density, as predicted by Nelson and Kosterlitz [92]. The location of the BKT transition — as deduced from the sharp disappearance of the superfluid fraction — is also consistent with the behaviour of the other physical quantities computed in chapter 3, including the decay of spatial correlation functions and the vortex unpairing.

Unlike explicitly constructing a dynamical simulation, our method is based on correlations that are calculated in the stationary thermal ensemble. This is a great practical advantage because it avoids the need for additional numerically costly simulations. It also avoids any concern that the dynamical perturbation in such simulations might disrupt the thermal background, resulting in a perturbed measurement of the superfluid fraction. Having said this, it would be interesting to compare our results to the superfluid fraction as deduced from the drag force felt by a moving impurity potential.

We finish by discussing some limitations and possible extensions to our technique. The main limitation is that the method appears to rely critically on translational symmetry; at the very least it is clear that the thought experiment used in the derivation doesn't make sense for a trapped system. This is unfortunate, since all experimental systems are necessarily trapped in some way or another and any quantitative comparisons with actual experiments must take this inhomogeneity into account. So far the superfluid fraction for the trapped system has been determined by using the universality result for the critical density in the homogeneous gas [106] in combination with the local density approximation [68, 13]. It would be interesting to be able to compute the superfluid fraction independently as we have done here, but it is not obvious how to generalise the derivation to this case.

In this thesis we have calculated the superfluid fraction only for the two-dimensional case. In 1D the tensor $\chi(\mathbf{k})$ is degenerate so there is no transverse component available, and the derivation fails to make sense. There is no such problem in 3D, and we expect our method to provide useful results in this case.

³Note that we do not expect an exact discontinuity in our calculated superfluid fraction due to the finite size of the system and the presence of statistical noise.

Effective 1D equations for trapped Bose-Einstein condensates

This chapter presents an ansatz for solutions to the 3D GPE in the quasi-1D regime. Our ansatz expresses the full 3D wavefunction in terms of a pair of 1D complex fields that describe the amplitude and width of an elongated BEC. We derive equations of motion for the 1D fields using the Lagrangian formalism and solve them numerically for several test scenarios. We compare with other 1D approaches, and the full 3D solution.

5.1 Introduction

The Gross-Pitaevskii equation (GPE) described in section 2.3 [Eq. (2.68)] is a remarkably successful tool for describing experimental BEC dynamics near $T = 0$. The reasons are simple: it is accurate for a large subset of experiments, and efficient numerical schemes are fairly simple to implement. Nevertheless, solving the GPE in three dimensions can be computationally demanding because of the size of the spatial grid required. In some cases it is possible to mitigate this problem by dimensional reduction; there are at least two possibilities: First, we may make use of symmetries in the physical situation to reduce the number of dimensions. For example, experimental systems commonly have cylindrically symmetric trapping potentials; if the initial state also has cylindrical symmetry we may then eliminate the angle variable, which reduces the simulation to two dimensions. Second, there are experimental situations where the system is strongly confined in one or two “transverse” dimensions so that the dynamics in those dimensions is particularly simple. We may then make an ansatz for the wavefunction that allows us to integrate out the transverse directions to produce a lower dimensional *effective equation* which may be solved on a much smaller numerical grid.

Elongated, cylindrically symmetric cigar-shaped BECs are routinely produced in the laboratory using a variety of trapping techniques including optical lattices [56], atom chips [64, 79, 44, 114] and other types of magnetic traps as first achieved by Görlitz *et al.* [54]. In parallel, several

studies (see, for example, [34, 120, 65, 20, 85]) have investigated the dynamics of dispersive shock waves¹ in BEC. Refs. [20, 85] combined both one dimensionality and shock wave generation by rapidly applying or removing a perturbing potential at the centre of a cigar-shaped BEC. In both of these studies the resulting shock waves were modelled numerically using a GPE or GPE-like mean-field approach. The calculations in Ref. [85] required a high spatial resolution in the longitudinal direction due to both the large size of the condensate and the need to resolve shock fronts at small length scales. As a result, these calculations were restricted to using a 1D effective equation for efficiency. The method chosen was the non-polynomial Schrödinger equation (NPSE) of Salasnich *et al.* [111], one of several methods that adiabatically eliminate the transverse motion and assume the shape of the transverse profile varies slowly as a function of the longitudinal coordinate.

While the NPSE is convenient, it is not obvious that the assumptions used in its derivation are valid when dealing with shocks. In particular, the existence of a shock implies rapid variation of the density and other system parameters with the longitudinal coordinate, and we would expect this to carry over to the shape of the transverse profile. In this chapter we relax both of the assumptions which go into deriving the NPSE — we include both the variation in the longitudinal direction and avoid making the adiabatic approximation.

An additional motivation for this work was to derive a 1D effective equation capable of simulating the expansion dynamics after turning off the trapping potential. These type of expansions are the standard experimental tool for imaging condensates, but simulating them directly is difficult due to the large size of the spatial grid required. The alternative 1D ansätze discussed in the next section eliminate the transverse velocity either implicitly or explicitly, and are therefore fundamentally unable to deal with expansion.

5.1.1 Previous work

We consider quasi-1D systems where the x and y coordinates correspond to the tightly confined transverse directions; the z coordinate is the weakly-trapped longitudinal direction. We assume that the transverse trapping potential is harmonic and cylindrically symmetric, while the longitudinal potential V_z is generic. The full potential is

$$V(\mathbf{x}) = V_z(z) + \frac{m\omega_{\perp}^2}{2}r^2, \quad (5.1)$$

¹A *dispersive* shock wave occurs when dispersion rather than dissipation dominates the physics at small length scales. Dispersive shock waves are characterised by pulse trains as seen in the simulations of Ref. [85], for example.

where $r = \sqrt{x^2 + y^2}$ is the radial coordinate, $\mathbf{x} = (x, y, z)$, m is the atomic mass and ω_\perp the angular frequency of the transverse harmonic potential.

In the quasi-1D regime, ω_\perp is large enough to prevent significant excitation of the transverse degrees of freedom. The simplest effective 1D equation may be derived by considering the case without interactions, $U_0 = 0$. In this case the full 3D wavefunction ϕ factorises and the transverse component has the functional form $\frac{1}{\sqrt{\pi}\sigma}e^{-r^2/2\sigma^2}$, where σ is the transverse width. In the perturbative regime² where U_0 is small but nonzero, this suggests the simple ansatz

$$\phi(\mathbf{x}, t) = \psi_1(z, t) \frac{1}{\sqrt{\pi}\sigma} e^{-r^2/2\sigma^2}, \quad (5.2)$$

with σ taken equal to the width of the non-interacting ground state. This ansatz works fairly well when U_0 is small and leads to the 1D GPE which describes the evolution of the unknown function ψ_1 :

$$i\hbar\partial_t\psi_1 = -\frac{\hbar^2}{2m}\partial_z^2\psi_1 + V\psi_1 + U_{1D}|\psi_1|^2\psi_1, \quad (5.3)$$

where U_{1D} is the effective 1D nonlinearity constant. For larger U_0 the transverse width increases substantially and the 1D GPE becomes rather inaccurate unless the value of σ is adjusted accordingly. An appropriate value may be computed using a variational calculation, assuming a constant density in z .

More sophisticated one-dimensional approximations have been investigated by several authors; we review those that are relevant to the current work below. An attempt has been made to keep to the notation used in the original papers, with some modifications for consistency.

The first of the more sophisticated approximations was the NPSE, as described by Salasnich *et al.* [111] in 2002. They used the ansatz

$$\phi(\mathbf{x}, t) = \frac{1}{\sqrt{\pi}\sigma(z, t)} e^{-r^2/2\sigma^2(z, t)} f(z, t), \quad (5.4)$$

along with the assumption that the width σ changes sufficiently slowly in the longitudinal direction that $\partial_z\sigma$ is negligible. The equation for f which results is known as the nonpolynomial Schrödinger equation due to the nonpolynomial nonlinear term that arises after eliminating the transverse width σ :

$$i\hbar\partial_t f = -\frac{\hbar^2}{2m}\partial_z^2 f + V_z f + \left[\frac{U_0}{2\pi a_\perp^2} \frac{|f|^2}{\sqrt{1+2a_s|f|^2}} + \frac{\hbar\omega_\perp}{2} \left(\frac{1}{\sqrt{1+2a_s|f|^2}} + \sqrt{1+2a_s|f|^2} \right) \right] f, \quad (5.5)$$

²The perturbative regime is defined by $a_s n_1 \ll 1$ where n_1 is the integrated 1D density — see, for example, Ref. [89].

where $a_{\perp} = \sqrt{\hbar/m\omega_{\perp}}$ is the transverse length scale. Implicit in the ansatz from Eq. (5.4) is the assumption that transverse dynamics given by changes in σ are much faster than the dynamics of f in which we are interested. To see this, note that the phase of ϕ does not depend on the coordinate r , which implies the absence of a radial superfluid velocity — the ansatz does not support transverse dynamics independently of the field f .

This lack of transverse dynamics was addressed more carefully by Muñoz Mateo and Delgado [89, 90], who showed that the general ansatz

$$\phi(\mathbf{x}, t) = \varphi(r, n_1(z, t)) \phi_1(z, t) \quad (5.6)$$

may be used to derive an equation for ϕ_1 via the adiabatic elimination of the transverse degrees of freedom. Here φ is some family of radial wavefunctions, parametrised by the local 1D condensate density n_1 . By assuming that $\partial_z n_1 = 0$, making the adiabatic approximation, and integrating out the transverse direction, they obtained the remarkably simple equation

$$i\hbar \partial_t \phi_1 = -\frac{\hbar^2}{2m} \partial_z^2 \phi_1 + V_z \phi_1 + \mu_{\perp}(n_1) \phi_1, \quad (5.7)$$

where

$$\mu_{\perp}(n_1) = \iint dx dy \varphi^* \left(-\frac{\hbar^2}{2m} \nabla_{\perp}^2 + \frac{m\omega_{\perp}^2}{2} r^2 + U_0 n_1 |\varphi|^2 \right) \varphi \quad (5.8)$$

is the local chemical potential. (We will refer to Eq. (5.7) as the Muñoz Mateo-Delgado equation (MDE) from now on.) Choosing the formula $\mu_{\perp}(n_1) = \hbar\omega_{\perp} \sqrt{1 + 4a_s n_1}$ to interpolate between known limits at large and small U_0 , they were able to obtain very accurate predictions of intrap oscillations with a range of smooth initial conditions, when benchmarked against a full 3D calculation [89].

When the transverse dynamics are important or the condensate width varies rapidly with the longitudinal coordinate z , we expect approaches based on the adiabatic approximation to become less accurate. Kamchatnov *et al.* [73] investigated a different generalisation of the NPSE ansatz,

$$\phi(\mathbf{x}, t) = \frac{1}{\sqrt{\pi}b(z, t)} e^{-r^2/2b^2(z, t)} e^{(i/2)\alpha(z, t)r^2} \tilde{\psi}(z, t), \quad (5.9)$$

using a variational formalism. The added generality and lack of assumptions regarding the z derivatives of the fields were motivated by the desire to describe solitons involving short length scales. The Kamchatnov ansatz results in coupled equations³ for the two real fields $b(z, t)$, $\alpha(z, t)$ and one complex field $\tilde{\psi}(z, t)$, which were used to analytically investigate solitonic solutions in

³We refer the reader to Ref. [73] for details of the equations — they are not used further here, and are reasonably complex to write down.

various limits, in addition to small amplitude linear waves. An important feature is the inclusion of the phase factor α , which allows for “breathing mode” dynamics in the transverse direction.

In this chapter we introduce a variation of the Kamchatnov approach, which — while mathematically equivalent — leads to more compact equations. In contrast to the earlier work, we investigate general numerical solutions. We show cases where the solutions are accurate, and highlight a number of generic difficulties which will arise for any similar factorisation ansatz.

5.2 Formalism

5.2.1 The Ansatz

We consider the following approximation to the 3D wavefunction:

$$\phi(\mathbf{x}, t) = \psi(z, t)e^{\chi(z, t)r^2}. \quad (5.10)$$

Here both ψ and χ are complex time-varying fields in one spatial dimension. The field χ is related to the transverse Gaussian width and transverse superfluid speed, respectively, via

$$\sigma(z) = 1/\sqrt{-2 \operatorname{Re} \chi(z)} \quad \text{and} \quad v_\perp(z, r) = 2|\operatorname{Im} \chi(z)|r. \quad (5.11)$$

The field ψ contains density and longitudinal phase information along with a normalisation factor for χ ; with this ansatz, the 1D density has the form

$$n(z, t) = \iint dx dy |\phi(\mathbf{x}, t)|^2 = \frac{\pi |\psi(z, t)|^2}{-[\chi(z, t) + \chi^*(z, t)]}. \quad (5.12)$$

In what follows we will often suppress time and space arguments to avoid excessive notational clutter.

5.2.2 Equations of motion

To derive the equations of motion we use a time dependent variational formalism essentially the same as that described in section 2.3.1. We begin by computing a Lagrangian for the one dimensional effective theory by integrating out the transverse dimensions. Taking derivatives of the effective Lagrangian and putting these into the Euler-Lagrange equations then yields the equations of motion in the usual manner.

One possible Lagrangian⁴ for the three-dimensional classical $|\phi|^4$ field theory is

$$\mathcal{L}[\phi] = \int d\mathbf{x} \frac{i\hbar}{2} (\phi^* \dot{\phi} - \dot{\phi}^* \phi) - \mathcal{H}[\phi], \quad (5.13)$$

where the Hamiltonian functional is given by

$$\mathcal{H}[\phi] = \int d\mathbf{x} \left(\frac{\hbar^2}{2m} |\nabla \phi|^2 + V |\phi|^2 + \frac{U_0}{2} |\phi|^4 \right). \quad (5.14)$$

Differentiating \mathcal{L} with respect to ϕ^* yields the 3D GPE via the Euler-Lagrange equations. If instead we substitute our ansatz from Eq. (5.10) into Eq. (5.13) we may perform the Gaussian integrals over the transverse coordinates, x and y . We then obtain a reduced Lagrangian which is a function of the fields ψ and χ :

$$\begin{aligned} \mathcal{L}[\psi, \chi] = \pi \int dz \left\{ \frac{i\hbar}{2} \left[\frac{\dot{\chi} - \dot{\chi}^*}{(\chi + \chi^*)^2} |\psi|^2 - \frac{\psi^* \dot{\psi} - \dot{\psi} \psi^*}{\chi + \chi^*} \right] \right. \\ \left. - \frac{\hbar^2}{2m} \left[\frac{4|\chi|^2 |\psi|^2}{(\chi + \chi^*)^2} - \frac{2|\partial_z \chi|^2 |\psi|^2}{(\chi + \chi^*)^3} - \frac{|\partial_z \psi|^2}{\chi + \chi^*} + \frac{\psi \partial_z \psi^* \partial_z \chi + \psi^* \partial_z \psi \partial_z \chi^*}{(\chi + \chi^*)^2} \right] \right. \\ \left. + \frac{V_z |\psi|^2}{\chi + \chi^*} - \frac{m\omega_\perp^2 |\psi|^2}{2(\chi + \chi^*)} + \frac{U_0 |\psi|^4}{4(\chi + \chi^*)} \right\}. \quad (5.15) \end{aligned}$$

As usual, stationarity of the action implies the Euler-Lagrange equations for our two fields, which are

$$\frac{d}{dt} \left(\frac{\delta \mathcal{L}}{\delta \dot{\chi}^*} \right) = \frac{\delta \mathcal{L}}{\delta \chi^*} \quad \text{and} \quad \frac{d}{dt} \left(\frac{\delta \mathcal{L}}{\delta \dot{\psi}^*} \right) = \frac{\delta \mathcal{L}}{\delta \psi^*}. \quad (5.16)$$

Performing the functional derivatives (see appendix A), rearranging and simplifying the results gives the equations of motion for χ and ψ :

$$\begin{aligned} i\hbar \dot{\chi} &= \frac{\hbar^2}{2m} \left[-\partial_z^2 \chi - 2 \frac{\partial_z \psi}{\psi} \partial_z \chi + 4 \frac{(\partial_z \chi)^2}{\chi + \chi^*} - 4\chi^2 \right] + \frac{m\omega_\perp^2}{2} + \frac{U_0}{4} |\psi|^2 (\chi + \chi^*), \\ i\hbar \dot{\psi} &= \frac{\hbar^2}{2m} \left[-\partial_z^2 \psi + \frac{2(\partial_z \chi)^2 \psi}{(\chi + \chi^*)^2} - 4\psi \chi \right] + V_z \psi + \frac{3U_0}{4} |\psi|^2 \psi. \end{aligned} \quad (5.17)$$

We highlight two relevant features of the equations above that may not be immediately obvious: First, Eqs. (5.17) are energy conserving due to the time independence of the Lagrangian⁵. Second, our choice of generalised coordinates in Eq. (5.10) forces us to give up the Hamiltonian structure of the phase space in exchange for a compact parametrisation of the wavefunction.

⁴This form of the Lagrangian is manifestly symmetrical with respect to ϕ and ϕ^* but we could equally well have chosen a form more similar to that used in section 2.3.1.

⁵The sophisticated way to say this is that the time translation symmetry of the Lagrangian implies energy conservation via Noether's theorem.

This is not physically problematic, but does affect our choice of numerical methods.

5.2.3 Conservation of normalisation and conserved current

An important property of any low-dimensional effective equation arising from the GPE is the conservation of normalisation of the wavefunction; this corresponds to conservation of the total number of atoms during time evolution. The total normalisation is

$$N = \int d\mathbf{x} |\phi(\mathbf{x})|^2 = \int dz \frac{\pi|\psi|^2}{-(\chi + \chi^*)}. \quad (5.18)$$

In principle, N is a function of time; to prove that it is not and the total number is conserved, we show that $\dot{N} = 0$. Expanding \dot{N} and inserting the equations of motion to remove the resulting factors of $\dot{\psi}$ and $\dot{\chi}$, we have:

$$\dot{N} = -\pi \int dz \left[\frac{\dot{\psi}\psi^* + \dot{\psi}^*\psi}{\chi + \chi^*} - \frac{|\psi|^2(\dot{\chi} + \dot{\chi}^*)}{(\chi + \chi^*)^2} \right] \quad (5.19)$$

$$= \frac{i\hbar\pi}{2m} \int dz \left[\left(-\frac{\psi^*\partial_z^2\psi}{\chi + \chi^*} + \frac{|\psi|^2\partial_z^2\chi}{(\chi + \chi^*)^2} + \frac{2\psi^*\partial_z\psi\partial_z\chi}{(\chi + \chi^*)^2} - \frac{2|\psi|^2(\partial_z\chi)^2}{(\chi + \chi^*)^3} \right) - \text{c.c.} \right] \quad (5.20)$$

where c.c. stands for the complex conjugate of the preceding bracketed term. Finally, we may remove all second order derivatives using integration by parts, assuming that the boundary terms are zero⁶. The resulting terms cancel out, implying that $\dot{N} = 0$ and normalisation is conserved.

Conservation of normalisation suggests that there should also be a conserved one-dimensional current $j(z, t)$, obeying the equation

$$\dot{n}(z, t) = -\partial_z j(z, t). \quad (5.21)$$

Computing \dot{n} directly from the definition in Eq. (5.12) and rearranging shows that the appropriate conserved current may be written

$$j = \frac{\hbar}{2mi} \left[\pi \left(\frac{|\psi|^2\partial_z\chi}{(\chi + \chi^*)^2} - \frac{\psi^*\partial_z\psi}{\chi + \chi^*} \right) - \text{c.c.} \right]. \quad (5.22)$$

If we define the 1D wavefunction $\eta = \psi\sqrt{-\pi/(\chi + \chi^*)}$ with the properties $n(z, t) = |\eta(z, t)|^2$

⁶In a periodic system the individual boundary terms are generally nonzero but nevertheless correctly cancel due to periodicity.

and $\arg[\phi(0, 0, z, t)] = \arg[\eta(z, t)]$, the current takes on a somewhat more familiar form

$$j = \frac{\hbar}{2mi} \left(\eta^* \partial_z \eta - \eta \partial_z \eta^* + |\eta|^2 \frac{\partial_z(\chi^* - \chi)}{\chi + \chi^*} \right). \quad (5.23)$$

The usual conserved probability current for quantum mechanics consists of the first two terms in the expression for j . The extra term is a consequence of integrating out the radial structure.

5.2.4 Ground states

We make use of ground states as physically reasonable initial conditions for the numerical simulations in the next section. Here we describe the ground state equations to be solved, along with the simplest approximate solution.

Ground states of Eqs. (5.17) have no transverse dynamics, that is, $\dot{\chi} = 0$. Further, the transverse speed v_\perp is zero which implies $\text{Im } \chi = 0$. The time evolution of ψ is the simple phase rotation $\dot{\psi} = (-i\mu/\hbar)\psi$ for some chemical potential μ , and we may assume for simplicity that ψ is real at $t = 0$. With these observations, we see that the ground state obeys the time-independent system of equations

$$\begin{aligned} 0 &= \frac{\hbar^2}{2m} \left[-\partial_z^2 \chi - 2 \frac{\partial_z \psi}{\psi} \partial_z \chi + 2 \frac{(\partial_z \chi)^2}{\chi} - 4\chi^2 \right] + \frac{A}{2} + \frac{U_0}{2} |\psi|^2 \chi, \\ 0 &= \frac{\hbar^2}{2m} \left[-\partial_z^2 \psi + \frac{(\partial_z \chi)^2 \psi}{2\chi^2} - 4\psi \chi \right] + (V_z - \mu)\psi + \frac{3U_0}{4} |\psi|^2 \psi, \end{aligned} \quad (5.24)$$

where both χ and ψ are real.

In cases where the local terms dominate in Eqs. (5.24) — slow spatial variation, large interactions, or large densities — we may ignore the z derivative terms. This is equivalent to using χ and ψ as determined from a spatially homogeneous system with the same local density, and is therefore known as a local density approximation (LDA).

Setting $\partial_z \psi = 0$ and $\partial_z \chi = 0$, Eqs. (5.24) reduce to

$$\begin{aligned} -\frac{2\hbar^2}{3m} \chi^2 + \frac{2}{3}(\mu - V)\chi + \frac{A}{2} &= 0, \\ |\psi|^2 &= \frac{4}{3U_0} \left[\frac{2\hbar^2}{m} \chi + (\mu - V) \right]. \end{aligned} \quad (5.25)$$

These algebraic equations are easily solved for χ and ψ . We take the branch of the square root such that $\chi < 0$ so that the transverse width in Eq. (5.11) is a real number. We note that Eqs. (5.25) are analogous to the well known Thomas-Fermi approximation for the GPE ground state.

5.3 Numerical Simulations

In this section we present numerical results for the simulation of Eqs. (5.17) along with comparisons to full 3D simulations and other proposed effective 1D equations. Our 1D equations are discretised on a uniform grid with periodic boundary conditions; derivatives are calculated spectrally via the Fourier transform. We use the standard fourth order Runge-Kutta integrator for time stepping. For the numerical simulations we use units such that $\hbar = m = 1$, and have arbitrarily chosen $\omega_{\perp} = 10$ and a periodic domain $-8 \leq z < 8$ with wavefunction normalisation $N = 1000$. With these choices, nonlinear behaviour becomes strongly apparent at values of the interaction strength $U_0 \sim 1$.

The full 3D system was assumed to be cylindrically symmetric, allowing a reduction to a two-dimensional equation. We simulated this 2D equation using the XMDS2 software [1] with a Bessel Fourier basis. Time stepping was achieved using an adaptive fourth/fifth order Runge-Kutta solver⁷ for time stepping.

5.3.1 Ground states

A common approach for finding the ground state of the GPE is to minimise the Hamiltonian via continuous-time steepest-descent optimisation, the so-called “imaginary time” method⁸. Unfortunately, this method does not always work well for the equations presented here, because the energy depends only weakly on the transverse degrees of freedom χ in regions of low density. This means that convergence of χ to the ground state value can be extremely slow in regions of high potential V_z .

To avoid this problem, we solve a discretised version of the ground state differential equations (5.24) directly using Newton’s method. The spatial discretisation and derivative calculation used here is the same as for the dynamical equations. The desired normalisation for the wavefunction is obtained by treating the chemical potential as one of the unknowns and adjoining the normalisation condition to the set of equations.

Convergence of Newton’s method for this system requires a starting point which is quite close to the true solution. For this we make use of the LDA solution Eq. (5.25). We find that the presence of low density regions where the LDA gives $\psi = 0$ results in Newton’s method failing with a singular Jacobian. A simple method to avoid this problem is to smooth the LDA solution with a spatial filter with exponentially decaying tails.

⁷named ARK45 in the XMDS software

⁸We note that for the equations presented here, steepest-descent minimisation is *not* equivalent to replacing the time t with $i\tau$ and evolving in τ . This is because the equations of motion are not Hamiltonian after the transformation in Eq. (5.10).

We also needed to find ground states for the NPSE, MDE and 3D equations for comparison purposes. The imaginary time method is sufficient in these cases because the equations arise directly from 3D or effective 1D Hamiltonians, albeit with sometimes unusual nonlinearity terms. To satisfy a given normalisation we use a slightly unusual version of the imaginary time method where the chemical potential μ is treated as an unknown and adjusted continuously along with the wavefunction (see appendix C for details).

5.3.2 Test case: soliton formation

As a test case, we consider evolution in a circular waveguide ($V_z = 0$) after releasing the ground state of the barrier potential $V_z = C \exp(-10z^2)$ at time $t = 0$. The barrier intensity $C = 50$ is chosen such that the density inside the barrier is depressed to few percent of the background density to induce strongly nonlinear evolution.

After finding the ground state, the barrier is turned off and the atoms fill the resulting hole. For sufficiently large interaction strength U_0 , the excitations are carried away as a combination of sound waves and a pair of grey solitons form, as shown in Fig. 5.1(a). The associated phase profile presented in Fig. 5.1(c) shows the expected jump in phase for a soliton across the density depression at $t = 1$. In addition there is a decrease in the width and associated radial flow.

To evaluate the 1D model we compare with the full 3D results, and also to two alternative effective 1D equations: the NPSE and the MDE. A comparison of the soliton speed as a function of interaction strength is shown in Fig. 5.2(b). We see that our equations and the MDE are competitive over most of the range, however the MDE wins out at larger interaction strengths. Of note is the NPSE results that are significantly less accurate than our equations, even though the same transverse Gaussian ansatz is used.

We note that the MDE has a particular advantage in the reproduction of accurate ground states: the interpolating form used for the local chemical potential accurately represents the important properties of the transverse wavefunction. In contrast, both methods which assume a transverse Gaussian profile are less accurate when the transverse wavefunction is deformed due to the interaction energy. Nevertheless, our equations are significantly more accurate than the NPSE, which indicates that derivatives of the transverse wavefunction parameters have a measurable effect for these initial conditions.

Further examination of Fig. 5.1(a) shows an apparent oscillation in the soliton depth coupled with the emission of sound waves. This feature is not present in the MDE or NPSE simulations but is present in the full 3D results. To quantify the effect, Fig. 5.2(a) shows the soliton depth as a function of time. We see that the initial condition causes an oscillation in the transverse width which feeds back into the soliton depth in this case. While the width cannot be observed

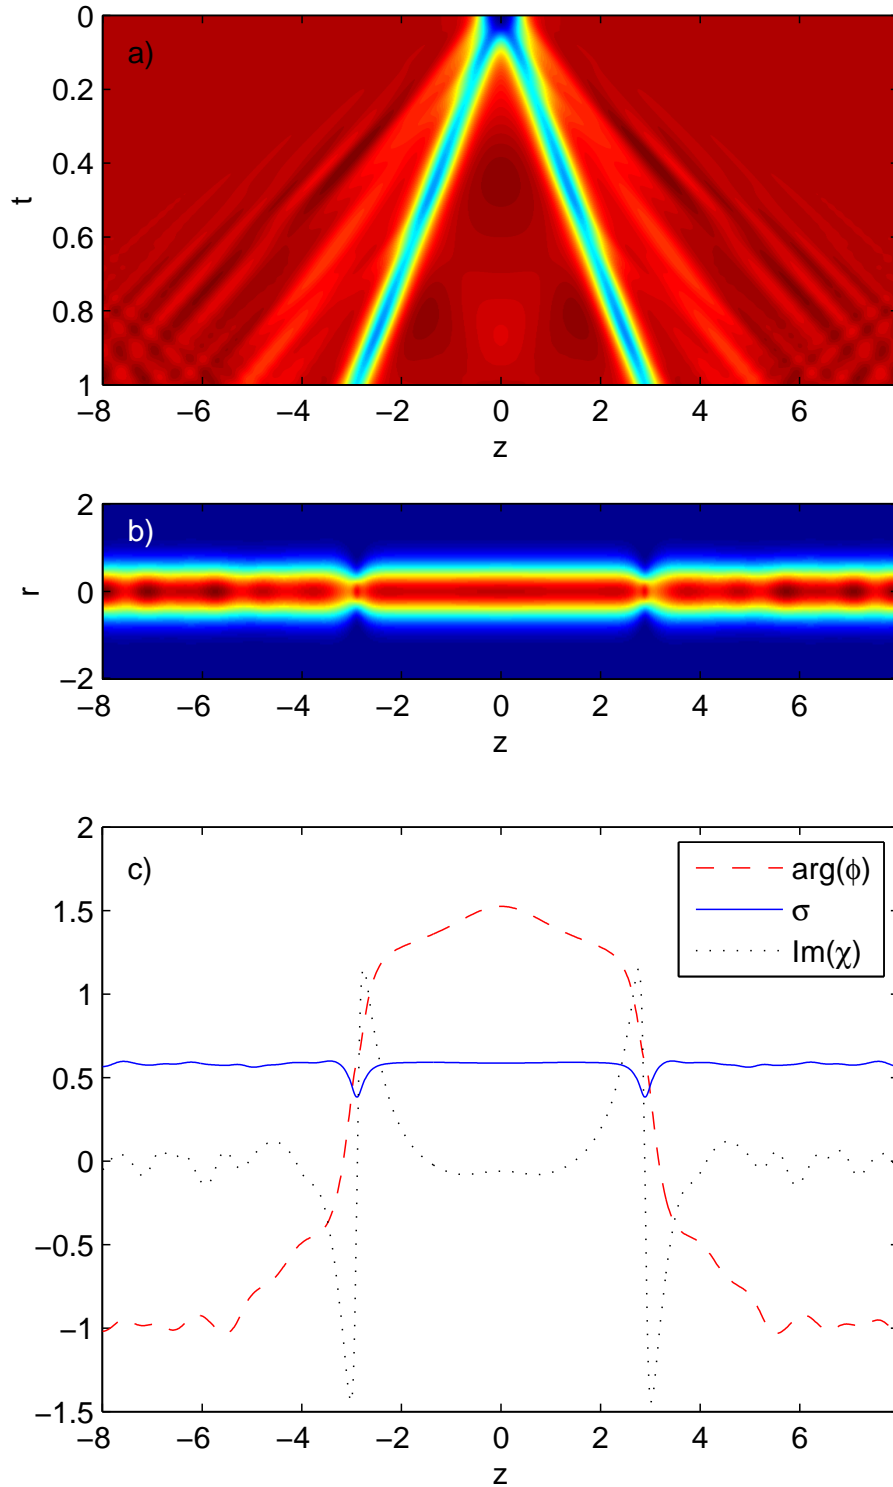


Figure 5.1: Solitons on a constant density background. The initial condition is the ground state of a Gaussian barrier potential. (a) Spatial density as a function of time. (b) Density slice through the centre of the inferred 3D cloud at $t = 1$. (c) Detail of the fields at $t = 1$ — dashes indicate the phase profile, $\arg(\psi)$, the solid line is the width of the Gaussian, σ , and the dotted line shows the scaling factor for the radial velocity, $\text{Im} \chi$.

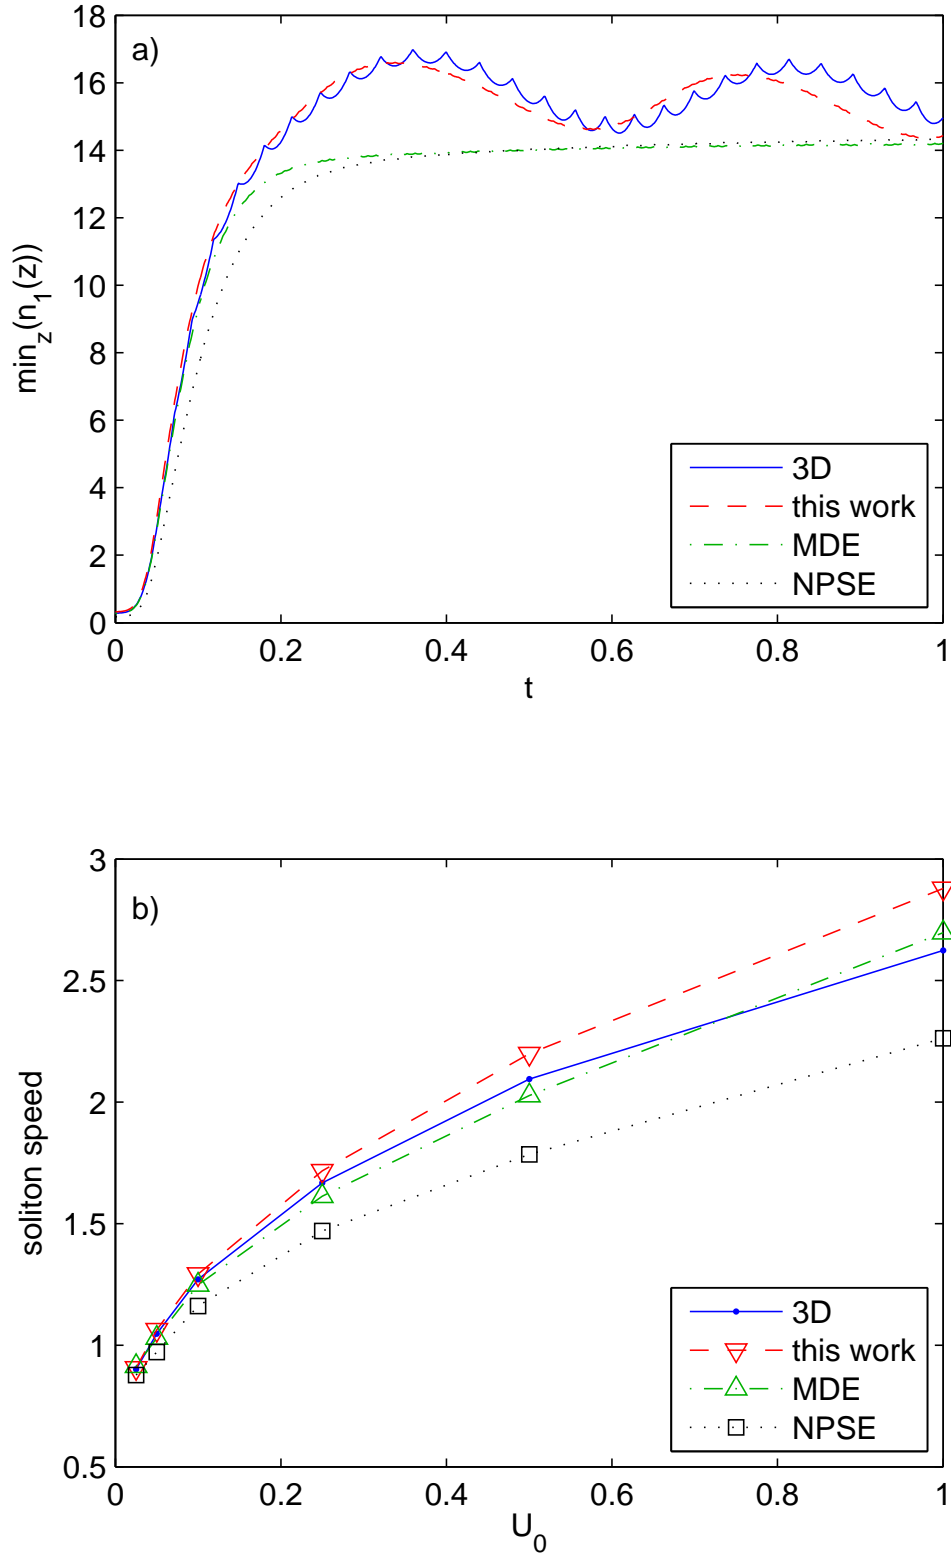


Figure 5.2: Comparison between various simulation methods for the soliton waveguide problem of section 5.3.2. (a) The soliton depth as a function of time for $U_0 = 0.25$. The high frequency oscillations in the 3D solution correspond to a small excitation of the higher order transverse modes. (b) The soliton speed as a function of the interaction strength U_0 .

after integrating out the transverse direction, the soliton depth can.

5.3.3 Transverse shock wave formation

We have presented a case above where our equations are competitive with other effective 1D equations for modelling the full 3D dynamics. Nevertheless, we have found recurring stability problems over the course of many numerical experiments with varying initial conditions. Such instability is typically characterised by sharp shock-like features that initially develop in the transverse width χ . We observed that the shocks are often associated with the density passing close to zero, and inspecting the equation of motion for χ in Eqs. (5.17) suggests that the term containing $1/\psi$ might be responsible: When ψ is very small this term becomes very large, potentially resulting in stiff equations. Explicit schemes such as the classic fourth order Runge-Kutta method become unstable in the presence of stiffness unless an unmanageably small time step is used.

We tested this hypothesis by making use of a Matlab implementation [38] of the RADAU5 algorithm [63]. RADAU5 is a fifth order instance of the Radau IIA class of implicit Runge-Kutta algorithms, and is known for its exceptional stability properties [63]. Nevertheless we found the results to be inconclusive: while it seemed to help for some initial conditions, there were certainly cases where the solutions diverged regardless of the type of numerical integrator used.

After further investigation, we believe we have identified the underlying cause of these problems: The ideal transverse width is not always a continuous function of the transverse slices of the 3D wavefunction. To make matters worse, discontinuities develop dynamically and it is not generally obvious which initial conditions will lead to problems.

To demonstrate the issue, we fix $U_0 = 0.1$ and consider initial conditions that are the ground state of the trapping potential

$$V(z) = V_0 \left[1 - e^{-10z^2} \right]. \quad (5.26)$$

By changing V_0 , the nonuniformity of the initial conditions can be increased from a completely homogeneous system ($V_0 = 0$) to one in which all the density is concentrated near $z = 0$. For sufficiently strong initial nonuniformity ($V_0 = 50$) we observe the formation of discontinuities in the transverse width of the 3D system at a finite time $t \approx 0.17$ after the pulse is released, as shown in Fig. 5.3. We computed the 3D width by nonlinear least squares fitting of a Gaussian to the 3D transverse density profile shown in Fig. 5.4. Also shown in Fig. 5.3 is the width arising from our ansatz, according to Eq. (5.11). While the width oscillations are not reproduced accurately, the existence and position of the discontinuity is correct. As a result, numerical methods that assume χ is smooth will generally fail within the next few time steps after the

snapshot shown in Fig. 5.3.




Figure 5.3: Transverse Gaussian width (standard deviation σ) of the wavefunction density from our ansatz, compared to a 3D simulation at time $t = 0.17$. This is a zoomed view in the z axis, clearly showing the discontinuities at $z \approx 1.6$.

The existence of discontinuities comes as a surprise; one might expect that continuous evolution of a smooth underlying field should lead to smooth fitting parameters. The 3D transverse profile in Fig. 5.4 shows that the underlying density is smooth as expected, and that the transverse profile is well behaved along most of the z coordinate. However, in the region of very low density near the discontinuity ($z \approx 1.6$), the transverse profile has two peaks as shown in Fig. 5.5. The discontinuity comes about because the least squares solution jumps suddenly from fitting the thin central peak to fitting a wider combination of the central and secondary peaks. Even though the density is low in the problematic region, the discontinuity in width is enough to destabilise the numerical solution.

We expect that *any* ansatz for the 3D wavefunction containing a single transverse width parameter will suffer from similar problems. It is somewhat ironic that the NPSE, MDE and similar approaches based on the adiabatic approximation remain stable and somewhat accurate even in regimes where our equations fail: Such methods avoid the problem by assuming that




Figure 5.4: Density snapshot from a 3D simulation showing radial structure at time $t = 0.17$, as in Fig. 5.3.

the transverse profile depends only weakly on z !

5.4 Conclusion

In this chapter we have examined the general Gaussian ansatz for the quasi-1D Bose gas tightly confined in a transverse 2D harmonic trap. We have compared the ansatz to two alternative 1D effective equations, the NPSE and MDE. We have found that some transient behaviours — such as an oscillation in soliton depth after release from a trap — can only be captured with an ansatz that allows transverse breathing motion. Such motion is prohibited by the NPSE and MDE due to an assumption that the transverse degrees of freedom adiabatically follow the local density. Further, we have shown that the method is more accurate than the NPSE due to the inclusion of derivatives of the transverse width.

However, we have found that our general Gaussian ansatz has several important failings that make it impractical for general use. First and most severely, discontinuities appear in the Gaussian width parameter for a wide range of sufficiently excited initial conditions. Second, finding ground states numerically is significantly more difficult than with competing 1D approximations. Third, the assumption of a Gaussian profile implies less accurate ground states than the MDE when the interaction strength is sufficiently strong.

The appearance of discontinuities in the width at finite time is an interesting issue which we trace to the appearance of two peaks in the transverse structure. Although this problem typically occurs in regions of low density and energy, it spells disaster for numerical solutions. We expect it is a generic feature of any similar ansatz which fits the transverse wavefunction using a single width parameter. One could consider additional parameters to get a better fit




Figure 5.5: Two transverse Gaussian fits (red) to the radial profile of the 3D wavefunction (blue) at adjacent grid points in the z direction. This example shows how a discontinuity in the fitting parameters may arise even when the underlying function to be fitted is continuous. The grid points in z span the position of the discontinuity in σ shown in Fig. 5.3.

to the transverse wavefunction, such as taking the next one or two harmonic oscillator basis states, scaled by a Gaussian width parameter. However, the resulting equations would clearly be more complex to derive — and more unwieldy to use — than those presented here.

Our results show that out of the three effective 1D equations considered, the MDE [89, 90] has the best combination of reliability, simplicity and relative accuracy.

Conclusion

In this thesis we have used numerical simulations to investigate the physics of ultracold Bose gases confined to one- and two-dimensional geometries.

Our two-dimensional simulations were carried out using the PGPE classical field technique. The PGPE allowed us to investigate the finite temperature regime surrounding the BKT superfluid phase transition, in a parameter regime that is relevant to current experiments. We directly computed the superfluid and condensate fractions and found that these drop to zero at the same temperature: in our finite-sized system there is no clear separation between BEC and BKT transitions. This is in contrast to what is expected based on the theory of the infinite-sized system, emphasising the importance of finite-size effects.

The unbinding of vortex pairs above the BKT transition is one of the defining features of the transition, and we developed a coarse-graining technique to quantify the number of unpaired vortices in the system. Results proved consistent with several other characteristic features, including the decay of the first-order correlation function and the sharp drop of the superfluid fraction to zero.

With a detailed microscopic theory at hand, we were able to simulate and validate the experimental measurement technique of Ref. [60], including the method used to deduce the first-order correlations from interference experiments. In addition, our simulations suggest that direct observation of vortex pairs is not possible due to limited imaging resolution. However, for this same reason it is likely that the dislocations observed in experimental interference patterns correspond to free vortices, and accordingly are a strong indicator of the BKT transition.

To investigate the BKT transition we required a method for calculating the superfluid fraction from a PGPE classical field simulation. Ideally such a method would be based on properties of the equilibrium thermal ensemble as ergodically sampled by the PGPE. We were able to derive such a method, using linear response theory to relate the superfluid fraction to the autocorrelations of the momentum density operator. Our method provided physically reasonable results that were consistent with the other physical properties of the 2D gas.

One-dimensional effective equations allow for efficient simulation of elongated cigar-shaped BECs. We derived equations of motion using the Lagrangian formalism with a Gaussian ansatz,

and solved these numerically for several test cases. We found that in certain cases our equations are more accurate than effective equations based on the adiabatic approximation. However, they are numerically unstable for a wide range of initial conditions. We investigated this instability and found that it is an inherent weakness of any similar ansatz which relies on a single variational width parameter.

In summary, this thesis contributes original insight into the physics of ultracold Bose gases in low dimensions, shedding light on recent experiments and supplying several new theoretical and numerical tools.

References

- [1] XMDS: eXtensible Multi-Dimensional Simulator, version 2.0, 2011. Software homepage at <http://www.xmds.org>.
- [2] A. Altland and B. Simons. *Condensed matter field theory*. Cambridge University Press, Cambridge, 2006.
- [3] B. P. Anderson, P. C. Haljan, C. A. Regal, D. L. Feder, L. A. Collins, C. W. Clark, and E. A. Cornell. Watching dark solitons decay into vortex rings in a Bose-Einstein condensate. *Phys. Rev. Lett.*, 86:2926, 2001.
- [4] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. Observation of Bose-Einstein condensation in a dilute atomic vapor. *Science*, 269(5221):198, 1995.
- [5] J. R. Anglin and W. Ketterle. Bose-Einstein condensation of atomic gases. *Nature*, 416:211, 2002.
- [6] K. J. Arrow and L. Hurwicz. Gradient method for concave programming, I Local results. In K. J. Arrow, L. Hurwicz, and H. Uzawa, editors, *Studies in Linear and Non-Linear Programming*. Stanford University Press, Stanford, California, 1958.
- [7] S. D. Bartlett, T. Rudolph, and R. W. Spekkens. Dialogue concerning two views on quantum coherence: factist and fictionist. *Int. J. Quantum Inf.*, 4:17, 2006.
- [8] G. Baym. The microscopic description of superfluidity. In R. C. Clark and G. H. Derrick, editors, *Mathematical Methods in Solid State and Superfluid Theory*. Oliver & Boyd, Edinburgh, 1969.
- [9] V. L. Berezinskii. Destruction of long-range order in one-dimensional and two-dimensional systems having a continuous symmetry group I. classical systems. *Sov. Phys. JETP*, 32:493, 1971.

-
- [10] A. Bezett, E. Toth, and P. B. Blakie. Two-point correlations of a trapped interacting Bose gas at finite temperature. *Phys. Rev. A*, 77(2):023602, 2008.
- [11] D. J. Bishop and J. D. Reppy. Study of the superfluid transition in two-dimensional ^4He films. *Phys. Rev. Lett.*, 40:1727, 1978.
- [12] R. N. Bisset and P. B. Blakie. Quantitative test of mean-field description of a trapped two-dimensional Bose gas. *Phys. Rev. A*, 80:045603, 2009.
- [13] R. N. Bisset, M. J. Davis, T. P. Simula, and P. B. Blakie. Quasicondensation and coherence in the quasi-two-dimensional trapped Bose gas. *Phys. Rev. A*, 79(3):033626, 2009.
- [14] P. B. Blakie, A. S. Bradley, M. J. Davis, R. J. Ballagh, and C. W. Gardiner. Dynamics and statistical mechanics of ultra-cold Bose gases using c-field techniques. *Adv. Phys.*, 57:363, 2008.
- [15] P. B. Blakie and M. J. Davis. Classical region of a trapped Bose gas. *J. Phys. B*, 40(11):2043–2053, 2007.
- [16] R. Bonifacio, L. M. Narducci, and E. Montaldi. Conditions for existence of a diagonal representation for quantum mechanical operators. *Phys. Rev. Lett.*, 16:1125, 1966.
- [17] S. N. Bose, O. Theimer, Translator, and Budh Ram, Translator. The beginning of quantum statistics: A translation of “Planck’s law and the light quantum hypothesis”. *Am. J. Phys.*, 44(11):1056, 1976. Translation of S. N. Bose. Plancks Gesetz und Lichtquantenhypothese, *Z. Phys.* 26(3):178–181 (1924).
- [18] E. G. Cavalcanti, C. J. Foster, M. D. Reid, and P. D. Drummond. Bell inequalities for continuous-variable correlations. *Phys. Rev. Lett.*, 99:210405, 2007.
- [19] P. M. Chaikin and T. C. Lubensky. *Principles of Condensed Matter Physics*. Cambridge University Press, Cambridge, 1995.
- [20] J. J. Chang, P. Engels, and M. A. Hoefer. Formation of dispersive shock waves by merging and splitting Bose-Einstein condensates. *Phys. Rev. Lett.*, 101(17):170404, 2008.
- [21] M. L. Chiofalo and M. P. Tosi. Output from Bose condensates in tunnel arrays: The role of mean-field interactions and of transverse confinement. *Phys. Lett. A*, 268(4-6):406, 2000.
- [22] R. Churchill. *Complex Variables and Applications*. McGraw-Hill, second edition, 1960.

-
- [23] P. Cladé, C. Ryu, A. Ramanathan, K. Helmerson, and W. D. Phillips. Observation of a 2D Bose gas: From thermal to quasicondensate to superfluid. *Phys. Rev. Lett.*, 102(17):170401, 2009.
- [24] F. Dalfovo, S. Giorgini, L. P. Pitaevskii, and S. Stringari. Theory of Bose-Einstein condensation in trapped gases. *Rev. Mod. Phys.*, 71:463, 1999.
- [25] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle. Bose-Einstein condensation in a gas of sodium atoms. *Phys. Rev. Lett.*, 75:3969, 1995.
- [26] M. J. Davis, R. J. Ballagh, and K. Burnett. Dynamics of thermal Bose fields in the classical limit. *J. Phys. B*, 34(22):4487–4512, 2001.
- [27] M. J. Davis and P. B. Blakie. Calculation of the microcanonical temperature for the classical Bose field. *J. Phys. A*, 38(48):10259–10271, 2005.
- [28] M. J. Davis and P. B. Blakie. Critical temperature of a trapped Bose gas: Comparison of theory and experiment. *Phys. Rev. Lett.*, 96(6):060404, 2006.
- [29] M. J. Davis and S. A. Morgan. Microcanonical temperature for a classical field: Application to Bose-Einstein condensation. *Phys. Rev. A*, 68(5):053615, 2003.
- [30] M. J. Davis, S. A. Morgan, and K. Burnett. Simulations of bose fields at finite temperature. *Phys. Rev. Lett.*, 87(16):160402, 2001.
- [31] M. J. Davis, S. A. Morgan, and K. Burnett. Simulations of thermal Bose fields in the classical limit. *Phys. Rev. A*, 66(5):053618, 2002.
- [32] B. DeMarco and D. S. Jin. Onset of Fermi degeneracy in a trapped atomic gas. *Science*, 285(5434):1703, 1999.
- [33] P. A. M. Dirac. On the theory of quantum mechanics. *Proc. R. Soc. Lond*, 112(762):661, 1926.
- [34] Z. Dutton, M. Budde, C. Slowe, and L. V. Hau. Observation of quantum shock waves created with ultra-compressed slow light pulses in a Bose-Einstein condensate. *Science*, 293(5530):663, 2001.
- [35] P. Ehrenfest and J. R. Oppenheimer. Note on the statistics of nuclei. *Phys. Rev.*, 37:333, 1931. Neutrons had not yet been discovered at the time this was published.

- [36] A. Einstein. Quantum theory of the monoatomic ideal gas. *Sitzungsber. Preuss. Akad. Wiss.*, page 261, 1924. English translation available online at <http://www.condmat.uni-oldenburg.de/TeachingSP/SP.html> (last accessed 26 November 2011).
- [37] A. Einstein. Quantum theory of the monoatomic ideal gas, second treatise. *Sitzungsber. Preuss. Akad. Wiss.*, page 3, 1925. English translation available online at <http://www.condmat.uni-oldenburg.de/TeachingSP/SP.html> (last accessed 26 November 2011).
- [38] Ch. Engstler. Matlab implementation of the Radau IIA method. <http://na.uni-tuebingen.de/projects.shtml> (last accessed 11 March 2010).
- [39] L. Erdős, B. Schlein, and H.-T. Yau. Rigorous derivation of the Gross-Pitaevskii equation. *Phys. Rev. Lett.*, 98:040404, 2007.
- [40] E. Fermi. On the quantization of the monoatomic ideal gas. *Rend. Lincei*, 3:145, 1926. English translation by A. Zannoni, arXiv:cond-mat/9912229v1.
- [41] Z. Ficek and M. R. Wahiddin. *Quantum Optics — Fundamentals and Applications*. International Islamic University Malaysia, Kuala Lumpur, 2004.
- [42] A. Filinov, N. V. Prokof'ev, and M. Bonitz. Berezinskii-Kosterlitz-Thouless transition in two-dimensional dipole systems. *Phys. Rev. Lett.*, 105:070401, 2010.
- [43] D. S. Fisher and P. C. Hohenberg. Dilute Bose gas in two dimensions. *Phys. Rev. B*, 37:4936, 1988.
- [44] R. Folman, P. Krüger, J. Schmiedmayer, J. Denschlag, and C. Henkel. Microscopic atom optics: from wires to an atom chip. *Adv. At., Mol., Opt. Phys.*, 48:263–356, 2002.
- [45] D. Forster. *Hydrodynamic Fluctuations, Broken Symmetry, and Correlation Functions*. Benjamin, Reading, Massachusetts, 1975.
- [46] C. J. Foster, P. B. Blakie, and M. J. Davis. Vortex pairing in two-dimensional Bose gases. *Phys. Rev. A*, 81:023623, 2010.
- [47] C. W. Gardiner and M. J. Davis. The stochastic Gross-Pitaevskii equation: II. *J. Phys. B*, 36(23):4731–4753, 2003.
- [48] C. W. Gardiner, P. Zoller, R. J. Ballagh, and M. J. Davis. Kinetics of Bose-Einstein condensation in a trap. *Phys. Rev. Lett.*, 79(10):1793–1796, 1997.

-
- [49] S. A. Gardiner and S. A. Morgan. Number-conserving approach to a minimal self-consistent treatment of condensate and noncondensate dynamics in a degenerate Bose gas. *Phys. Rev. A*, 75:043621, 2007.
- [50] C. Gies and D. A. W. Hutchinson. Coherence properties of the two-dimensional Bose-Einstein condensate. *Phys. Rev. A*, 70:043606, 2004.
- [51] L. Giorgetti, I. Carusotto, and Y. Castin. Semiclassical field method for the equilibrium Bose gas and application to thermal vortices in two dimensions. *Phys. Rev. A*, 76(1):013613, 2007.
- [52] R. J. Glauber. Coherent and incoherent states of the radiation field. *Phys. Rev.*, 131:2766, 1963.
- [53] H. Goldstein. *Classical Mechanics*. Addison-Wesley, second edition, 1980.
- [54] A. Görlitz, J. M. Vogels, A. E. Leanhardt, C. Raman, T. L. Gustavson, J. R. Abo-Shaeer, A. P. Chikkatur, S. Gupta, S. Inouye, T. Rosenband, and W. Ketterle. Realization of Bose-Einstein condensates in lower dimensions. *Phys. Rev. Lett.*, 87(13):130402, 2001.
- [55] I. S. Gradshteyn and I. M. Ryzhik. *Table of Integrals, Series, and Products*. Academic Press, San Diego, sixth edition, 2000.
- [56] M. Greiner, I. Bloch, O. Mandel, T. W. Hansch, and T. Esslinger. Exploring phase coherence in a 2D lattice of Bose-Einstein condensates. *Phys. Rev. Lett.*, 87(16), 2001.
- [57] W. Greiner and J. Reinhardt. *Field quantization*. Springer Verlag, 1996.
- [58] A. Griffin. Conserving and gapless approximations for an inhomogeneous Bose gas at finite temperatures. *Phys. Rev. B*, 53(14):9341, 1996.
- [59] Z. Hadzibabic. Private communication.
- [60] Z. Hadzibabic, P. Krüger, M. Cheneau, B. Battelier, and J. Dalibard. Berezinskii-Kosterlitz-Thouless crossover in a trapped atomic gas. *Nature*, 441(7097):1118–1121, 2006.
- [61] Z. Hadzibabic, S. Stock, B. Battelier, V. Bretin, and J. Dalibard. Interference of an array of independent Bose-Einstein condensates. *Phys. Rev. Lett.*, 93(18):180403, 2004.
- [62] S. A. Haine and M. T. Johnsson. Dynamic scheme for generating number squeezing in Bose-Einstein condensates through nonlinear interactions. *Phys. Rev. A*, 80:023611, 2009.

-
- [63] E. Hairer and G. Wanner. *Solving Ordinary Differential Equations II: Stiff and Differential-Algebraic Problems*. Springer Series in Comput. Math. Springer, Berlin, 2nd revised edition, 1996.
- [64] W. Hansel, P. Hommelhoff, T. W. Hansch, and J. Reichel. Bose-Einstein condensation on a microelectronic chip. *Nature*, 413(6855):498–501, 2001.
- [65] M. A. Hoefer, M. J. Ablowitz, I. Coddington, E. A. Cornell, P. Engels, and V. Schweikhard. Dispersive and classical shock waves in Bose-Einstein condensates and gas dynamics. *Phys. Rev. A*, 74:023623, 2006.
- [66] P. C. Hohenberg. Existence of long-range order in one and two dimensions. *Physical Review*, 158:383–386, 1967.
- [67] M. Holzmann, M. Chevallier, and W. Krauth. Universal correlations and coherence in quasi-two-dimensional trapped Bose gases. *arXiv preprint*, arXiv:0911.1704v1, 2009.
- [68] M. Holzmann and W. Krauth. Kosterlitz-Thouless transition of the quasi-two-dimensional trapped Bose gas. *Phys. Rev. Lett.*, 100(19):190402, 2008.
- [69] R. F. Hoskins. *Delta functions: an introduction to generalised functions*. Horwood Series in Mathematics & Applications. Horwood Publishing Ltd., Chichester, England, 1999.
- [70] R. Jackiw and A. Kerman. Time-dependent variational principle and the effective action. *Phys. Lett. A*, 71(2–3):158, 1979.
- [71] Y. Kagan and B. V. Svistunov. Evolution of correlation properties and appearance of broken symmetry in the process of Bose-Einstein condensation. *Phys. Rev. Lett.*, 79(18):3331, 1997.
- [72] Y. Kagan, B. V. Svistunov, and G. V. Shlyapnikov. Influence on inelastic processes of the phase transition in a weakly collisional two-dimensional Bose gas. *Sov. Phys. JETP*, 66:314, 1987.
- [73] A. M. Kamchatnov and V. S. Shchesnovich. Dynamics of Bose-Einstein condensates in cigar-shaped traps. *Phys. Rev. A*, 70(2):023604, 2004.
- [74] V. A. Kashurnikov, N. V. Prokof'ev, and B. V. Svistunov. Critical temperature shift in weakly interacting Bose gas. *Phys. Rev. Lett.*, 87:120402, 2001.
- [75] J. M. Kosterlitz and D. J. Thouless. Ordering, metastability and phase transitions in two-dimensional systems. *J. Phys. C*, 6:1181, 1973.

-
- [76] K. Kreutz-Delgado. The complex gradient operator and the CR-calculus. *arXiv preprint*, arXiv:0906.4835v1, 2009.
- [77] P. Krüger, Z. Hadzibabic, and J. Dalibard. Critical point of an interacting two-dimensional atomic Bose gas. *Phys. Rev. Lett.*, 99(4):040402, 2007.
- [78] L. Landau. Theory of the superfluidity of helium II. *Phys. Rev.*, 60:356, 1941.
- [79] A. E. Leanhardt, A. P. Chikkatur, D. Kielpinski, Y. Shin, T. L. Gustavson, W. Ketterle, and D. E. Pritchard. Propagation of Bose-Einstein condensates in a magnetic waveguide. *Phys. Rev. Lett.*, 89(4), 2002.
- [80] J. Leinaas and J. Myrheim. On the theory of identical particles. *Il Nuovo Cimento B*, 37:1, 1977.
- [81] C. Lubich. *From quantum to classical molecular dynamics: reduced models and numerical analysis*. European Mathematical Society, Zürich, 2008.
- [82] D. G. Luenberger. *Linear and Nonlinear Programming*. Addison-Wesley, Reading, Massachusetts, 1984.
- [83] L. Mathey and A. Polkovnikov. Supercritical superfluid and vortex unbinding following a quantum quench. *Phys. Rev. A*, 80(4):041601, 2009.
- [84] P. E. M. F. Mendonça, R. d. J. Napolitano, M. A. Marchioli, C. J. Foster, and Y.-C. Liang. Alternative fidelity measure between quantum states. *Phys. Rev. A*, 78:052330, 2008.
- [85] R. Meppelink, S. B. Koller, J. M. Vogels, P. van der Straten, E. D. van Ooijen, N. R. Heckenberg, H. Rubinsztein-Dunlop, S. A. Haine, and M. J. Davis. Observation of shock waves in a large Bose-Einstein condensate. *Phys. Rev. A*, 80:043606, 2009.
- [86] N. D. Mermin and H. Wagner. Absence of ferromagnetism or antiferromagnetism in one- or two-dimensional isotropic Heisenberg models. *Phys. Rev. Lett.*, 17:1133–1136, 1966.
- [87] K. Mølmer. Optical coherence: A convenient fiction. *Phys. Rev. A*, 55:3195, 1997.
- [88] S. A. Morgan. *A Gapless Theory of Bose-Einstein Condensation in Dilute Gases at Finite Temperature*. PhD dissertation, University of Oxford, 1999.
- [89] A. Muñoz Mateo and V. Delgado. Effective mean-field equations for cigar-shaped and disk-shaped Bose-Einstein condensates. *Phys. Rev. A*, 77(1):013617, 2008.

-
- [90] A. Muñoz Mateo and V. Delgado. Effective one-dimensional dynamics of elongated Bose-Einstein condensates. *Ann. Phys.*, 324(3):709, 2009.
- [91] M. Naraschewski and R. J. Glauber. Spatial coherence and density correlations of trapped Bose gases. *Phys. Rev. A*, 59(6):4595–4607, 1999.
- [92] D. R. Nelson and J. M. Kosterlitz. Universal jump in the superfluid density of two-dimensional superfluids. *Phys. Rev. Lett.*, 39(19):1201, 1977.
- [93] M. A. Nielsen and I. L. Chuang. *Quantum computation and quantum information*. Cambridge University Press, Cambridge; New York, 2000.
- [94] A. Papadopoulos. *Metric Spaces, Convexity and Nonpositive Curvature*. IRMA lectures in mathematics and theoretical physics; 6. European Mathematical Society, Zurich, 2005.
- [95] O. Penrose and L. Onsager. Bose-Einstein condensation and liquid helium. *Phys. Rev.*, 104:576, 1956.
- [96] C. J. Pethick and H. Smith. *Bose-Einstein Condensation in Dilute Gases*. Cambridge University Press, Cambridge, second edition, 2008.
- [97] D. S. Petrov, M. Holzmann, and G. V. Shlyapnikov. Bose-Einstein condensation in quasi-2D trapped gases. *Phys. Rev. Lett.*, 84(12):2551, 2000.
- [98] W. D. Phillips. Nobel lecture: Laser cooling and trapping of neutral atoms. *Rev. Mod. Phys.*, 70:721, 1998.
- [99] P. Pickl. Derivation of the time dependent Gross Pitaevskii equation with external fields. *arXiv preprint*, arXiv:1001.4894v2, 2010.
- [100] L. Pitaevskii and S. Stringari. *Bose-Einstein Condensation*. Oxford University Press, Oxford, 2003.
- [101] A. Polkovnikov. Phase space representation of quantum dynamics. *Annals of Physics*, 325(8):1790, 2010.
- [102] A. Polkovnikov, E. Altman, and E. Demler. Interference between independent fluctuating condensates. *Proc. Natl. Acad. Sci. U.S.A.*, 103(16):6125–6129, 2006.
- [103] V. N. Popov. *Functional Integrals in Quantum Field Theory and Statistical Physics*. Reidel, Dordrecht, 1983.

-
- [104] A. Posazhennikova. Weakly interacting, dilute Bose gases in 2D. *Rev. Mod. Phys.*, 78(4):1111, 2006.
- [105] N. Prokof'ev, O. Ruebenacker, and B. Svistunov. Critical point of a weakly interaction two dimensional Bose gas. *Phys. Rev. Lett.*, 87(27):270402, 2001.
- [106] N. Prokof'ev and B. Svistunov. Two-dimensional weakly interaction Bose gas in the fluctuation region. *Phys. Rev. A*, 66:043608, 2002.
- [107] N. P. Proukakis and B. Jackson. Finite-temperature models of Bose-Einstein condensation. *J. Phys. B*, 41(20):203002, 2008.
- [108] S.J. Putterman and P.H. Roberts. Classical non-linear waves in dispersive nonlocal media, and the theory of superfluidity. *Physica A*, 117:369, 1983.
- [109] R. Remmert. *Theory of Complex Functions*. Springer-Verlag, 1991.
- [110] H. H. Rugh. Dynamical approach to temperature. *Phys. Rev. Lett.*, 78(5):772–774, 1997.
- [111] L. Salasnich, A. Parola, and L. Reatto. Effective wave equations for the dynamics of cigar-shaped and disk-shaped Bose condensates. *Phys. Rev. A*, 65(4 B):0436141, 2002.
- [112] H. Salman, N. G. Berloff, and P. H. Roberts. From classical fields to two-fluid model of superfluidity: Emergent kinetics and local gauge transformations. In N.P. Proukakis, S.A. Gardiner, M.J. Davis, and M.H. Szymanska, editors, *Non-Equilibrium Quantum Gases at Finite Temperatures*. Imperial College Press, London, in press.
- [113] T. Sato, T. Suzuki, and N. Kawashima. Finite-temperature transition in a quasi-2D Bose gas trapped in the harmonic potential. *J. Phys. Conf. Ser.*, 150:032094, 2009.
- [114] S. Schneider, A. Kasper, Ch. vom Hagen, M. Bartenstein, B. Engeser, T. Schumm, I. Bar-Joseph, R. Folman, L. Feenstra, and J. Schmiedmayer. Bose-Einstein condensation in a simple microtrap. *Phys. Rev. A*, 67:023612, 2003.
- [115] D. V. Schroeder. *An introduction to thermal physics*. Addison Wesley, 2000.
- [116] D. Schumayer and D. A. W. Hutchinson. Thermodynamically activated vortex-dipole formation in a two-dimensional Bose-Einstein condensate. *Phys. Rev. A*, 75:015601, 2007.
- [117] V. Schweikhard, S. Tung, and E. A. Cornell. Vortex proliferation in the Berezinskii-Kosterlitz-Thouless regime on a two-dimensional lattice of Bose-Einstein condensates. *Phys. Rev. Lett.*, 99(3):030401, 2007.

-
- [118] T. P. Simula and P. B. Blakie. Thermal activation of vortex-antivortex pairs in quasi-two-dimensional Bose-Einstein condensates. *Phys. Rev. Lett.*, 96(2):020404, 2006.
- [119] T. P. Simula, M. J. Davis, and P. B. Blakie. Superfluidity of an interacting trapped quasi-two-dimensional Bose gas. *Phys. Rev. A*, 77(2):023618, 2008.
- [120] T. P. Simula, P. Engels, I. Coddington, V. Schweikhard, E. A. Cornell, and R. J. Ballagh. Observations on sound propagation in rapidly rotating Bose-Einstein condensates. *Phys. Rev. Lett.*, 94:080404, 2005.
- [121] S. Stock, Z. Hadzibabic, B. Battelier, M. Cheneau, and J. Dalibard. Observation of phase defects in quasi-two-dimensional Bose-Einstein condensates. *Phys. Rev. Lett.*, 95(19):190403, 2005.
- [122] P. Teller. *An Interpretive Introduction to Quantum Field Theory*. Princeton University Press, Princeton, New Jersey, 1995.
- [123] T. M. Wright, N. P. Proukakis, and M. J. Davis. Many-body physics in the classical-field description of a degenerate Bose gas. *Phys. Rev. A*, 84:023608, 2011.
- [124] C. N. Yang. Concept of off-diagonal long-range order and the quantum phases of liquid He and of superconductors. *Rev. Mod. Phys.*, 34:694, 1962.
- [125] S. M. Yoo, J. Ruostekoski, and J. Javanainen. Interference of two Bose-Einstein condensates. *J. Mod. Opt.*, 44(10):1763–1774, 1997.
- [126] V. I. Yukalov and H. Kleinert. Gapless Hartree-Fock-Bogoliubov approximation for Bose gases. *Phys. Rev. A*, 73:063612, 2006.
- [127] Ł. Zawitkowski, M. Gajda, and K. Rzażewski. Dynamics of a relative superflow between a Bose-Einstein condensate and the thermal cloud. *Phys. Rev. A*, 74:043601, 2006.

Mathematical techniques for classical field theory

A.1 The Wirtinger calculus: derivatives of nonholomorphic functions

The Wirtinger derivative is the suitable generalisation of the real derivative to general complex valued functions for the purposes of making linear approximations. As we will see below, the usual derivative of complex analysis is much too strict for these uses.

The derivative in complex analysis is defined using the familiar limit definition used in the real calculus, only with complex variables. Consider a function $f: \mathbb{C} \rightarrow \mathbb{C}$ and some point $z_0 \in \mathbb{C}$; the derivative of f at z_0 is

$$\left. \frac{df}{dz} \right|_{z_0} \equiv \lim_{z \rightarrow z_0} \frac{f(z) - f(z_0)}{z - z_0} \quad (\text{A.1})$$

(see, for example, [22, page 30]). If the limit converges to the same value irrespective of the way in which $z \rightarrow z_0$ we say that f is differentiable at z_0 . A function is called *holomorphic* on some domain D if it is differentiable at every point in an open set containing D . The existence conditions for complex derivatives defined like this are *much* stricter than for the real derivative. In particular, a necessary condition for the existence of the limit in Eq. (A.1) is the Cauchy Riemann conditions: If $f(x + iy) = u(x, y) + iv(x, y)$ with $u, v: \mathbb{R}^2 \rightarrow \mathbb{R}$, the Cauchy-Riemann conditions are $\frac{\partial u}{\partial x} = \frac{\partial v}{\partial y}$, and $\frac{\partial v}{\partial x} = -\frac{\partial u}{\partial y}$.

Deriving the equations of motion for a classical field theory from an action principle depends on taking derivatives of the action functional with respect to the fields and setting these to zero. That is, for an action $S[\psi]$ we would like to do something like set

$$\frac{\delta S}{\delta \psi} = 0, \quad (\text{A.2})$$

which should give us the equations of motion. However, the action S is generally a *real* function for any complex value of the field ψ , and as such cannot possibly satisfy the Cauchy-Riemann conditions¹. From the point of view of traditional complex analysis this immediately leads to a problem: how do we make use of the derivative $\delta S/\delta\psi$ when this derivative isn't even defined according to (some functional generalisation of) the definition in Eq. (A.1)? This is the main problem that we tackle in this section.

One complicating factor for our particular case is that the derivatives of interest are functional derivatives with respect to the field ψ . Luckily, the central problem remains intact if we forget about the fields and consider a nonholomorphic function f of a single complex variable z . This simplification is made throughout the rest of the section.

If we can't use the usual complex derivative for common problems, what is the alternative? Derivatives are all about linear approximations², so to understand what to do about the peculiar strictness of the Cauchy-Riemann conditions we consider a linear expansion of f about the point $z_0 = x_0 + iy_0$ in terms of the real variables x and y . Assuming that u and v are differentiable, this is simply

$$f(x + iy) \approx f(z_0) + \left(\frac{\partial u}{\partial x} \Big|_{z_0} + i \frac{\partial v}{\partial x} \Big|_{z_0} \right) \Delta x + \left(\frac{\partial u}{\partial y} \Big|_{z_0} + i \frac{\partial v}{\partial y} \Big|_{z_0} \right) \Delta y \quad (\text{A.3})$$

$$= f(z_0) + \frac{\partial f}{\partial x} \Big|_{z_0} \Delta x + \frac{\partial f}{\partial y} \Big|_{z_0} \Delta y \quad (\text{A.4})$$

where $\Delta x = x - x_0$ and $\Delta y = y - y_0$. This is a straightforward expansion using the real calculus so it's clear that this linear approximation is valid regardless of the Cauchy Riemann conditions. Using $\Delta z = z - z_0$, $\Delta z^* = z^* - z_0^*$ for the conjugate and suppressing the subscripts for brevity (all derivatives being evaluated at z_0), we have

$$f(z) \approx f(z_0) + \frac{\partial f}{\partial x} \Delta x + \frac{\partial f}{\partial y} \Delta y \quad (\text{A.5})$$

$$= f(z_0) + \frac{\partial f}{\partial x} \frac{1}{2} (\Delta z + \Delta z^*) + \frac{\partial f}{\partial y} \frac{1}{2i} (\Delta z - \Delta z^*) \quad (\text{A.6})$$

$$= f(z_0) + \frac{1}{2} \left(\frac{\partial f}{\partial x} - i \frac{\partial f}{\partial y} \right) \Delta z + \frac{1}{2} \left(\frac{\partial f}{\partial x} + i \frac{\partial f}{\partial y} \right) \Delta z^*. \quad (\text{A.7})$$

This is an elementary but interesting result: the general linear approximation to any real-differentiable complex-valued function may be expressed as a linear combination of Δz and

¹At least not in a nontrivial way: for a real function f we have $v(x, y) = 0$ and the derivative conditions immediately show that $u(x, y)$ must be a constant.

²In fact, the Fréchet derivative — an elegant generalisation of the derivative to Banach spaces — is by definition exactly the linear part of the affine approximation to a function.

Δz^* with the coefficients given by a simple combination of the real derivatives. Furthermore, we can now express the ratio inside the limit in Eq. (A.1) as

$$\frac{f(z) - f(z_0)}{\Delta z} = \frac{1}{2} \left(\frac{\partial f}{\partial x} - i \frac{\partial f}{\partial y} \right) + \frac{1}{2} \left(\frac{\partial f}{\partial x} + i \frac{\partial f}{\partial y} \right) e^{-2i\theta} \quad (\text{A.8})$$

where $\theta = \arg \Delta z$ which reveals the general form of the directional dependence of the limit. If f is to be complex differentiable, the value of the right hand side must be independent of z which implies $\frac{\partial f}{\partial x} + i \frac{\partial f}{\partial y} = 0$; rearranging we see that this is exactly an expression of the Cauchy Riemann conditions.

Coming back to Eq. (A.7), we *define* two new derivative operators

$$\frac{\partial f}{\partial z} \equiv \frac{1}{2} \left(\frac{\partial f}{\partial x} - i \frac{\partial f}{\partial y} \right) \quad \text{and} \quad \frac{\partial f}{\partial z^*} \equiv \frac{1}{2} \left(\frac{\partial f}{\partial x} + i \frac{\partial f}{\partial y} \right), \quad (\text{A.9})$$

which allow us to write the linear approximation as

$$f(z) \approx f(z_0) + \frac{\partial f}{\partial z} \Delta z + \frac{\partial f}{\partial z^*} \Delta z^*. \quad (\text{A.10})$$

At first sight this expression may seem like a mere curiosity that has been manufactured to resemble the real calculus of functions of two variables. On the contrary, it turns out that there is a powerful and general calculus for the operators $\frac{\partial}{\partial z}$ and $\frac{\partial}{\partial z^*}$. Most importantly, given a function $f(z)$ written in terms of z and z^* , it is possible to compute derivatives directly with the new operators using the familiar rules of calculus for two *real* variables, and *without* decomposing f into real and imaginary parts. This calculus is known as the Wirtinger calculus [109] after the Austrian mathematician Wilhelm Wirtinger³.

A.1.1 Properties of the Wirtinger derivative

We now discuss the general properties of the Wirtinger derivative as discussed in Ref. [109, §1.4]. Several properties are immediate from the definitions in Eq. (A.9):

- The derivative operators are linear, due to the linearity of the real partial derivatives.
- They obey the usual product rule.
- The behaviour under complex conjugation is given by

$$\left(\frac{\partial f}{\partial z} \right)^* = \frac{\partial f^*}{\partial z^*}. \quad (\text{A.11})$$

³It has also been called the $\mathbb{C}\mathbb{R}$ -calculus by Kreutz-Delgado [76].

- The derivatives satisfy the “independence of z and z^* ” property,

$$\frac{\partial z^*}{\partial z} = 0 \quad \text{and} \quad \frac{\partial z}{\partial z^*} = 0. \quad (\text{A.12})$$

The last property has great importance from a practical computational viewpoint, and we return to it after discussing the chain rule.

Using the conjugation property and the linear approximation formula in Eq. (A.10) it is straightforward to show that the Wirtinger derivatives obey a familiar looking chain rule. We have

$$f(g(z)) \approx f\left(g(z_0) + \frac{\partial g}{\partial z}\Delta z + \frac{\partial g}{\partial z^*}\Delta z^*\right) \quad (\text{A.13})$$

$$\approx f(g(z_0)) + \left(\frac{\partial f}{\partial g}\frac{\partial g}{\partial z} + \frac{\partial f}{\partial g^*}\frac{\partial g^*}{\partial z}\right)\Delta z + \left(\frac{\partial f}{\partial g}\frac{\partial g}{\partial z^*} + \frac{\partial f}{\partial g^*}\frac{\partial g^*}{\partial z^*}\right)\Delta z^* \quad (\text{A.14})$$

where $\frac{\partial f}{\partial g} \equiv \frac{\partial f}{\partial z}\big|_{g(z)}$ and we have used the linear approximation formula twice, first for g and then for f . Identifying the coefficient in front of the Δz using the linear approximation formula, we have shown that the chain rule for Wirtinger derivatives is

$$\frac{\partial}{\partial z}f(g(z)) = \frac{\partial f}{\partial g}\frac{\partial g}{\partial z} + \frac{\partial f}{\partial g^*}\frac{\partial g^*}{\partial z}, \quad (\text{A.15})$$

with the analogous expression for $\frac{\partial}{\partial z^*}f(g(z))$. It is worth noting the close resemblance of this chain rule with the formula from the more familiar calculus on \mathbb{R}^2 ,

$$\frac{\partial}{\partial x}f(g(x, y)) = \frac{\partial f}{\partial g_1}\frac{\partial g_1}{\partial x} + \frac{\partial f}{\partial g_2}\frac{\partial g_2}{\partial x} \quad (\text{A.16})$$

where $f, g: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and $g(x, y) = [g_1(x, y), g_2(x, y)]$.

The relationship with the complex derivative in Eq. (A.1) of a holomorphic function is simply that

$$\frac{df}{dz} = \frac{\partial f}{\partial z}; \quad (\text{A.17})$$

the definitions coincide for this important special case, as expected from the linear approximation formula. For the same reason, it's possible to see that the Cauchy-Riemann conditions have a particularly nice form in this formalism:

$$\frac{\partial f}{\partial z^*} = 0. \quad (\text{A.18})$$

The generalisation of the Wirtinger calculus to higher dimensions is completely straightforward from a calculational viewpoint, because the variables z_k are genuinely independent for distinct k . We summarise this here by asserting that everything works just as in the real multivariable calculus, with the main difference being that variables always occur in conjugate pairs.

A.1.2 Computing with Wirtinger derivatives

The independence property $\frac{\partial z^*}{\partial z} = 0$ — along with the other basic properties — leads to the commonly repeated⁴ mnemonic device

To compute with the Wirtinger derivatives, treat z and z^ as “independent variables” and proceed exactly as for the calculus of two real variables.*

This “independence” is clearly problematic in a purely algebraic sense, but is extremely useful for the purposes of computing derivatives because it results in all the right rules when interpreted correctly. It is often useful to write functions $f(z)$ redundantly as $f(z, z^*)$ to aid in applying the mnemonic. For example, $f(g(z))$ would become $f(g(z, z^*), g^*(z, z^*))$ and the chain rule follows after remembering the version for two independent real variables.

We now present two examples that illustrate how computing with the Wirtinger calculus avoids the need to break a function into real and imaginary parts. A simple but useful example is the derivative of the absolute value:

$$\frac{\partial}{\partial z}|z| = \frac{\partial}{\partial z}(zz^*)^{1/2} \quad (\text{A.19})$$

$$= \frac{1}{2}(zz^*)^{-1/2} \frac{\partial}{\partial z}(zz^*) \quad (\text{A.20})$$

$$= \frac{z^*}{2|z|} \quad (\text{A.21})$$

where the second line is an application of the chain rule with the square root as the outer function and noting that the second term that appears in the chain rule vanishes because $\frac{\partial}{\partial z^*} z^{1/2} = 0$. In a similar way, $\frac{\partial}{\partial z^*}|z| = \frac{z}{2|z|}$.

As a more complicated example, suppose we wanted to find the stationary points of the real valued function

$$f(z) = |(z + 1)^{10} + z^*|. \quad (\text{A.22})$$

(We present this example because it is closely related to finding the stationary points of an action which is generally also a real valued function of complex arguments.) For general complex-

⁴See, for example, Ref. [53, page 598].

valued functions, the stationary points are the points at which $\frac{\partial f}{\partial z} = 0$ and $\frac{\partial f}{\partial z^*} = 0$ simultaneously, but for real valued functions we have $\frac{\partial f}{\partial z^*} = \left(\frac{\partial f}{\partial z}\right)^*$ and it suffices to satisfy the first condition. Computing the derivative is an exercise in applying the chain rule to $h(g(z))$ with functions $h(z) = |z|$ and $g(z) = (z + 1)^{10} + z^*$:

$$\frac{\partial f}{\partial z} = \frac{\partial h}{\partial g} \frac{\partial g}{\partial z} + \frac{\partial h}{\partial g^*} \frac{\partial g^*}{\partial z} \quad (\text{A.23})$$

$$= \left(\frac{[(z + 1)^{10} + z^*]^*}{2|(z + 1)^{10} + z^*|} \right) (10(z + 1)^9) + \left(\frac{(z + 1)^{10} + z^*}{2|(z + 1)^{10} + z^*|} \right) (1) \quad (\text{A.24})$$

$$= \frac{10[(z^* + 1)^{10} + z](z + 1)^9 + (z + 1)^{10} + z^*}{2|(z + 1)^{10} + z^*|}. \quad (\text{A.25})$$

Therefore, finding the stationary points in this case involves solving the high order polynomial

$$10[(z^* + 1)^{10} + z](z + 1)^9 + (z + 1)^{10} + z^* = 0. \quad (\text{A.26})$$

This could also have been derived by finding and setting *both* real derivatives to zero, but doing so using the formalism of the Wirtinger calculus is more convenient in this case.

A.2 Functional derivatives

In the Lagrangian formulation of dynamics the action S is a *functional* — that is, a function taking other functions as arguments⁵. Stationarity of the action means that the derivatives with respect to the fields are zero; this section investigates what we mean by such derivatives and how to calculate them.

Throughout the section we make two simplifying assumptions to streamline the presentation: First, the functionals under consideration and their arguments are real. Second, the argument fields are functions of a single independent variable so that all integrals are one-dimensional. Lifting these assumptions presents no particular difficulty and is covered very briefly in the next section.

A.2.1 An example

As in the previous section, we start by looking at linear approximations. Consider for example the simple functional

$$S[f] \equiv \int dx \, g(x) f^2(x), \quad (\text{A.27})$$

⁵This is obviously true for field theories, but remains true when we have a finite number of dynamical variables because such variables are functions of time.

where g is some fixed function, and f may be anything that is well behaved enough so that the integral is finite. Evaluating S at $f(x) + \eta(x)$ for some small test function η and expanding around f , we have

$$S[f + \eta] = \int dx g(x) [f(x) + \eta(x)]^2 \quad (\text{A.28})$$

$$= \int dx g(x) [f^2(x) + 2f(x)\eta(x) + \eta^2(x)] \quad (\text{A.29})$$

$$\approx \int dx g(x) f^2(x) + \int dx 2g(x) f(x) \eta(x), \quad (\text{A.30})$$

where we have discarded the quadratic terms in η . We see that the change in S which results from changing f to $f + \eta$ is approximately a linear functional:

$$S[f + \eta] - S[f] \approx L_f[\eta] \equiv \int dx 2g(x) f(x) \eta(x). \quad (\text{A.31})$$

In some sense the linear functional L defined here actually *is* the derivative of S — known as the Fréchet derivative — but a slightly different definition is usually favoured by physicists, based on an analogy with the gradient.

To motivate the definition of functional derivative commonly used in physics we compare to the discrete setting. For a finite number of variables the analogue of S would be some function $R(\mathbf{f}) \equiv \sum_{i=1}^N g_i f_i^2$ for $\mathbf{f}, \mathbf{g} \in \mathbb{R}^N$. The linear approximation to R about \mathbf{f} is

$$R(\mathbf{f} + \mathbf{e}) - R(\mathbf{f}) \approx \sum_{i=1}^N 2g_i f_i e_i = \sum_{i=1}^N \frac{\partial R}{\partial f_i} e_i = \nabla R \cdot \mathbf{e}. \quad (\text{A.32})$$

The correspondences with the continuous case are $f \leftrightarrow \mathbf{f}$, $\eta \leftrightarrow \mathbf{e}$, $g \leftrightarrow \mathbf{g}$ and $x \leftrightarrow i$. With this in mind, it's clear that the function $2g \cdot f$ in Eq. (A.31) is acting in an analogous role to the gradient $\nabla R = [\frac{\partial R}{\partial f_1}, \dots, \frac{\partial R}{\partial f_N}] = [2g_1 f_1, \dots, 2g_N f_N]$. The *functional derivative* is therefore written using similar notation to the partial derivative:

$$\frac{\delta S}{\delta f(x)} = 2g(x) f(x). \quad (\text{A.33})$$

A.2.2 General definition

Roughly speaking, the *functional derivative* of a functional S evaluated at f is a distribution⁶ $\frac{\delta S}{\delta f}$ such that

$$S[f + \eta] - S[f] = \int dx \frac{\delta S}{\delta f(x)} \eta(x) + O(\|\eta\|^2) \quad (\text{A.34})$$

for all test functions η . (Note that a more formal and general definition may be given in terms of the Fréchet derivative⁷.) As indicated in the previous section, the distribution $\frac{\delta S}{\delta f}$ is a generalisation of the gradient vector from a finite to an infinite number of independent variables. In the case that the distribution is simply a function, we may evaluate it at x to give the *number* $\frac{\delta S}{\delta f(x)}$ which is intuitively the answer to the question “how much does S change if we change f by a small amount at the point x ”.

This intuition leads to the “physicist’s definition” of functional derivative:

$$\frac{\delta S}{\delta f(x)} = \lim_{\epsilon \rightarrow 0} \frac{S[f + \epsilon \delta_x] - S[f]}{\epsilon} \quad (\text{dubious!}) \quad (\text{A.35})$$

where $\delta_x(y) = \delta(x - y)$. The notion of poking f exactly at x using a delta spike is intuitive but this definition doesn’t make much mathematical — or even calculational — sense if taken literally. For example, consider attempting to calculate the derivative of the functional $\int dx g(x) f^2(x)$ from the previous section. Trying to compute $S[f + \epsilon \delta_x]$ then results in a term containing the integral of δ_x^2 which has no well-defined value. It is possible to repair these problems by considering a sequence $\{\Delta_{x,n}\}$ of nascent delta functions such that $\Delta_{x,n} \rightarrow \delta_x$ as $n \rightarrow \infty$, and taking the limit $\epsilon \rightarrow 0$ *before* the limit $n \rightarrow \infty$. However, it seems better to simply regard Eq. (A.35) as an aid to the intuition, and fall back to the idea of linear approximations for computational purposes.

A.2.3 Calculating with functional derivatives

Functional derivatives may usually be calculated directly by considering the linear expansion of the functional, but in practice this can be tedious. Therefore it is useful to derive a few general formulae that can be applied to a wide range of functionals. Here we present a few such formulae that are relevant to this thesis, along with an indication of how to derive them. We

⁶A *distribution* is also known as a generalised function and includes such objects as the Dirac delta “function”; see, for example, Ref. [69].

⁷Borrowing the definition of the Fréchet derivative, we might define $\frac{\delta S}{\delta f}$ as the distribution that satisfies

$$\lim_{\|\eta\| \rightarrow 0} \frac{1}{\|\eta\|} \left\| S[f + \eta] - S[f] - \int dx \frac{\delta S}{\delta f(x)} \eta(x) \right\| = 0.$$

direct the reader to Ref. [57, §2.3] for additional exposition.

The functionals of interest are generally of the form

$$S[f] = \int dy \, g(f(y), f'(y)) \quad (\text{A.36})$$

for some differentiable function g of two variables. The task is to create a linear approximation for

$$S[f + \eta] = \int dy \, g(f(y) + \eta(y), f'(y) + \eta'(y)). \quad (\text{A.37})$$

Writing $g^{(1,0)}$ and $g^{(0,1)}$ for the first derivatives with respect to the first and second arguments respectively, and briefly dropping the explicit y dependence of all functions for succinctness, the linear approximation is

$$S[f + \eta] = \int dy \, \left[g(f, f') + g^{(1,0)}(f, f')\eta + g^{(0,1)}(f, f')\eta' \right] + O(\eta^2) \quad (\text{A.38})$$

$$\approx S[f] + \int dy \, g^{(1,0)}(f, f')\eta + \int dy \, g^{(0,1)}(f, f')\eta'. \quad (\text{A.39})$$

This is almost in the right form, but the last term has the derivative η' of the test function rather than the test function itself. The derivative may be removed by integrating by parts and assuming that the boundary terms are zero. This is true for a periodic domain, an infinite domain where the fields and their derivatives decay at infinity, or any domain where the test function η is constrained to be zero at the boundary. Integrating by parts leads to

$$S[f + \eta] \approx S[f] + \int dy \, \left[g^{(1,0)}(f(y), f'(y)) - \frac{\partial}{\partial y} [g^{(0,1)}(f(y), f'(y))] \right] \eta(y) \quad (\text{A.40})$$

and as a result,

$$\frac{\delta S}{\delta f(x)} = g^{(1,0)}(f(x), f'(x)) - \frac{\partial}{\partial x} [g^{(0,1)}(f(x), f'(x))]. \quad (\text{A.41})$$

It is interesting to note that this is almost an expression of the Euler-Lagrange equation for a single dynamical variable. To see this, note that the derivative may more succinctly — though less explicitly — be written as

$$\frac{\delta S}{\delta f} = \frac{\partial g}{\partial f} - \frac{\partial}{\partial x} \left[\frac{\partial g}{\partial f'} \right] \quad (\text{A.42})$$

where by $\frac{\partial g}{\partial f}$ and $\frac{\partial g}{\partial f'}$ we mean the derivatives of g with respect to the first and second arguments respectively, evaluated at $(f(x), f'(x))$. Making the suggestive replacements $g \rightarrow L$, $x \rightarrow t$,

Table A.1: Some useful functional derivatives. The notations $g^{(1,0)}$ and $g^{(0,1)}$ mean derivatives of g with respect to its first and second arguments respectively.

Functional S	Functional derivative $\frac{\delta S}{\delta f(x)}$
$\int dy g(f(y))$	$g'(f(x))$
$\int dy g(y)f'(y)$	$-g'(x)f(x)$
$\int dy g(f(y), f'(y))$	$g^{(1,0)}(f(x), f'(x)) - \frac{\partial}{\partial x} [g^{(0,1)}(f(x), f'(x))]$
$f(y)$	$\delta(x - y)$

$f \rightarrow q$, and $q' \equiv \dot{q}$, we have

$$\frac{\delta S}{\delta q} = \frac{\partial L}{\partial q} - \frac{\partial}{\partial t} \left[\frac{\partial L}{\partial \dot{q}} \right], \quad (\text{A.43})$$

and setting $\frac{\delta S}{\delta q} = 0$ yields the Euler-Lagrange equation in its usual form.

Table A.1 presents a summary of several useful functional derivatives, most arising as special cases of Eq. (A.41).

Details of two-dimensional simulations

B.1 Simulation using the PGPE

Here we outline our procedure for determining the properties of the **C** region and the steps used to create initial states for the PGPE solver. The **C** region itself is characterised by the cutoff momentum K , while the initial states are characterised by the energy $E_{\mathbf{C}}$ and number $N_{\mathbf{C}}$. We want to obtain values of these three properties that are consistent with a specified temperature T and total number of atoms N .

B.1.1 Hartree-Fock-Bogoliubov analysis

To generate an initial estimate of the **C** region parameters we solve the self-consistent Hartree-Fock-Bogoliubov (HFB) equations in the so-called Popov approximation¹ [58] to find an approximate thermal state for the system at a temperature T . The resulting state is a Bose Einstein distribution of quasiparticles interacting only via the mean-field, expressed in terms of the quasiparticle amplitudes $u_{\mathbf{k}}$ and $v_{\mathbf{k}}$.

Occupations for the **C** region field may be computed directly from the quasiparticle occupations via

$$n_{\mathbf{k}} = (u_{\mathbf{k}}^2 + v_{\mathbf{k}}^2) N_B(E_{\mathbf{k}}) + v_{\mathbf{k}}^2, \quad (\text{B.1})$$

where N_B is the Bose-Einstein distribution and $E_{\mathbf{k}}$ is the quasiparticle energy that is obtained by solving the Bogoliubov-de Gennes equations self-consistently [58]. This allows us to compute the cutoff as the maximum value of $\|\mathbf{k}\|$ consistent with sufficient modal occupation:

$$K = \max\{\|\mathbf{k}\| : n_{\mathbf{k}} \geq n_{\text{cut}}\}. \quad (\text{B.2})$$

We choose $n_{\text{cut}} = 5$ for the sufficient occupation condition on the **C** region modes.

¹It is amusing to note the commentary in Ref. [126] that maintains that Popov himself never suggested this trick, and would only have deemed it valid very near the transition temperature!

The number of atoms below the cutoff may be computed directly from the sum of the condensate number N_0 and the number of **C** region excited state atoms, $N_{1\mathbf{C}}$:

$$N_{\mathbf{C}} = N_0 + N_{1\mathbf{C}}, \quad \text{where} \quad N_{1\mathbf{C}} = \sum_{\mathbf{k} \in \mathbf{C} \setminus \{\mathbf{0}\}} n_{\mathbf{k}}. \quad (\text{B.3})$$

For the total energy below the cutoff, we use the expression

$$E_{\mathbf{C}} = \frac{\hbar^2}{mL^2} \left(\frac{gN_0^2}{2} + \lambda N_{1\mathbf{C}} - gN_{1\mathbf{C}}^2 \right) + \sum_{\mathbf{k} \in \mathbf{C} \setminus \{\mathbf{0}\}} E_{\mathbf{k}} [N_B(E_{\mathbf{k}}) - v_{\mathbf{k}}^2] \quad (\text{B.4})$$

where $\lambda = g(N_0 + 2N_{1\mathbf{C}})$. Rearranging, this is

$$E_{\mathbf{C}} = \frac{\hbar^2 g}{2mL^2} (N_{\mathbf{C}}^2 + N_{1\mathbf{C}}^2) + \sum_{\mathbf{k} \in \mathbf{C} \setminus \{\mathbf{0}\}} E_{\mathbf{k}} [N_B(E_{\mathbf{k}}) - v_{\mathbf{k}}^2]. \quad (\text{B.5})$$

The expression in Eq. (B.5) differs from Eq. (22) of Ref. [58] as we have retained the zeroth order (constant) terms that are required to match the energy scale of the HFB analysis to the zero point of energy in the classical field simulations.

B.1.2 Initial conditions for fixed total number

A simple comparison between simulations at varying temperatures can only be carried out if the total number of atoms is fixed. This presents a problem in our simulations: although the number of atoms and energy of the **C** region can be directly specified (see section B.1.3), we may only determine the total number after performing a simulation. This is because the number of atoms in the **I** region depends on the temperature and chemical potential that are calculated by ergodic averaging of the **C** region simulations.

Formally, this may be stated as a root finding problem: solve

$$N(N_{\mathbf{C}}, E_{\mathbf{C}}) = N_{\text{tot}} \quad (\text{B.6})$$

with initial guess provided by the solution to the HFB analysis in section B.1.1. Although both $N_{\mathbf{C}}$ and $E_{\mathbf{C}}$ affect the total number N , we choose to fix $N_{\mathbf{C}}$ to the initial guess and to vary $E_{\mathbf{C}}$ until the desired total number is found.

We note that evaluating the function $N(N_{\mathbf{C}}, E_{\mathbf{C}})$ is very computationally expensive and difficult to fully automate because it involves a simulation and several steps of analysis. For this reason we use a nonstandard root finding procedure: For the first iteration we simulate three energies about the initial guess $E_{\mathbf{C}}$ such that the results crudely span N_{tot} ; these three

simulations can be performed in parallel which significantly reduces the time to a solution. A second guess was obtained by quadratic fitting of $E_{\mathbf{C}}$ as a function of N which gives N accurate to within about 5% of N_{tot} . An additional iteration using the same interpolation method takes N to within 0.3%, which we consider sufficient.

We note that changing $E_{\mathbf{C}}$ during the root finding procedure means we have no direct control over the final temperature of each specific simulation. In our case this is not a problem because we only require a range of temperatures spanning the transition. In principle one could solve for a given temperature by allowing $N_{\mathbf{C}}$ to vary in addition to $E_{\mathbf{C}}$.

B.1.3 Initial conditions for given $E_{\mathbf{C}}$ and $N_{\mathbf{C}}$

We compute initial conditions for the \mathbf{C} region field in a similar way to Ref. [31]. Using the representation for the \mathbf{C} region given by Eq. (3.9), the task is to choose appropriate values for the $\{c_{\mathbf{n}}\}$. As a first approximation, choose the smallest value for a momentum cutoff K' such that the field with coefficients

$$c_{\mathbf{n}} = \begin{cases} Ae^{i\theta_{\mathbf{n}}} & \text{for } 0 < \|\mathbf{k}\| \leq K', \\ 0 & \text{for } |\mathbf{k}| > K', \end{cases} \quad (\text{B.7})$$

has energy greater than $E_{\mathbf{C}}$. Here A is chosen so that the field has normalisation corresponding to $N_{\mathbf{C}}$ atoms, and $\theta_{\mathbf{n}}$ is a randomly chosen phase which is fixed for each mode at the start of the procedure. The random phases allow us to generate many unique random initial states at the same energy.

By definition, the field defined by Eq. (B.7) has energy slightly above the desired energy. This problem is solved by mixing it with the lowest energy state:

$$c_{\mathbf{n}} = \begin{cases} A'e^{i\theta_0} & \text{for } \mathbf{n} = \mathbf{0}, \\ 0 & \text{elsewhere,} \end{cases} \quad (\text{B.8})$$

using a root finding procedure to converge on the desired energy $E_{\mathbf{C}}$. The scheme generates random realisations of a non-equilibrium field with given $E_{\mathbf{C}}$ and $N_{\mathbf{C}}$ which are then simulated to equilibrium before using ergodic averaging for computing statistics.

B.2 I region integrals

Our assumed self-consistent Wigner function (section 3.3.2) for the **I** region atoms takes a particularly simple form in the homogeneous case:

$$W(\mathbf{k}, \mathbf{x}) = \frac{1}{(2\pi)^2} \frac{1}{e^{(\hbar^2 \mathbf{k}^2 / 2m + 2\hbar^2 g n_{\mathbf{C}} / m - \mu_{\mathbf{C}}) / k_B T} - 1}. \quad (\text{B.9})$$

The above-cutoff density may then be found by direct integration:

$$n_{\mathbf{I}}(\mathbf{x}) = \int_{\|\mathbf{k}\| \geq K} d^2 \mathbf{k} W_{\mathbf{I}}(\mathbf{k}, \mathbf{x}), \quad (\text{B.10})$$

$$= -\frac{1}{\lambda_{\text{dB}}^2} \ln \left[1 - e^{-(\hbar^2 K^2 / 2m + 2\hbar^2 g n_{\mathbf{C}} / m - \mu_{\mathbf{C}}) / k_B T} \right]. \quad (\text{B.11})$$

In a similar way, the assumed Wigner function allows any desired physical quantity to be estimated via a suitable integral. A particular quantity of interest in the current work is the first-order correlation function, which can be obtained from the Wigner function as [91]

$$G_{\mathbf{I}}^{(1)}(\mathbf{x}, \mathbf{x}') = \int_{\|\mathbf{k}\| \geq K} d^2 \mathbf{k} e^{-i\mathbf{k} \cdot (\mathbf{x} - \mathbf{x}')} W_{\mathbf{I}}(\mathbf{k}, \frac{\mathbf{x} + \mathbf{x}'}{2}). \quad (\text{B.12})$$

This integral is of the general form

$$I_1(\mathbf{r}) \equiv \int_{\|\mathbf{k}\| > K} d^2 \mathbf{k} \frac{e^{-i\mathbf{k} \cdot \mathbf{r}}}{e^{A\mathbf{k}^2 + B} - 1}, \quad (\text{B.13})$$

for constants A and B . Noting that I_1 depends only on the length r of $\|\mathbf{r}\|$, and transforming k to polar coordinates (κ, θ) , we have

$$I_1(\mathbf{r}) = \int_K^\infty d\kappa \frac{\kappa}{e^{A\kappa^2 + B} - 1} \int_0^{2\pi} d\theta e^{-ir\kappa \cos \theta}, \quad (\text{B.14})$$

$$= \int_K^\infty d\kappa \frac{\kappa}{e^{A\kappa^2 + B} - 1} 2 \left[\Gamma\left(\frac{1}{2}\right) \right]^2 J_0(r\kappa), \quad (\text{B.15})$$

(see Ref. [55, p902] for the Bessel function identity).

Thus we obtain $G_{\mathbf{I}}^{(1)}(\mathbf{x}, \mathbf{x}')$ in terms of a one-dimensional integral which may be performed numerically:

$$G_{\mathbf{I}}^{(1)}(\mathbf{x}, \mathbf{x}') = \frac{1}{2\pi} \int_K^\infty d\kappa \frac{\kappa J_0(\kappa \|\mathbf{x} - \mathbf{x}'\|)}{e^{(\hbar^2 \kappa^2 / 2m + 2\hbar^2 g n_{\mathbf{C}} / m - \mu_{\mathbf{C}}) / k_B T} - 1}. \quad (\text{B.16})$$

B.3 Vortex detection

The defining feature of a “charge- m ” vortex is that the phase θ of the complex field $\psi(\mathbf{x}) = |\psi(\mathbf{x})|e^{i\theta(\mathbf{x})}$ changes continuously from 0 to $2m\pi$ around any closed curve that circles the vortex core. We express our field ψ on a discrete grid in position space; the aim of vortex detection is then to determine which grid plaquettes (that is, sets of four adjacent grid points) contain vortex cores.

To obtain the phase winding about a plaquette, first consider the phase at two neighbouring grid points A and B. We are interested in the unwrapped phase difference $\Delta\theta_{AB}$ between the grid points; unwrapping ensures that the phase is *continuous* between A and B. (In the discrete setting such continuity is poorly defined; the best we can do is to correct for the possibility of 2π phase jumps by adding or subtracting factors of 2π so that $|\Delta\theta_{AB}| < \pi$.) The unwrapped phase differences around a grid plaquette tell us a total phase change $\theta_{\text{wrap}} = \sum_i \Delta\theta_{i,i+1} = 2m\pi$ where $m \in \mathbb{Z}$ is the winding number or “topological charge”.

Due to the necessity of unwrapping the phase, a four-point grid plaquette cannot unambiguously support vortices with charge larger than one. Luckily, such vortices are energetically unfavourable in 2D Bose gases [100, page 83] so we need only concern ourselves with detecting vortices with winding number ± 1 in this work. The positions obtained from a given run of our vortex detection algorithm are the labelled $\{\mathbf{r}_i^+\}$ and $\{\mathbf{r}_i^-\}$ for winding numbers $+1$ and -1 , respectively.

An imaginary time method with fixed normalisation

Ground states for the GPE are often found using the so-called “imaginary time” method [21]. The prescription is to replace the time t with the “imaginary time” $\tau = it$ and the potential V with $V - \mu$ in the equation of motion. One then selects a desired chemical potential μ , and evolves the new equations in τ until a stationary state ϕ_0 is reached. This stationary state corresponds to a ground state of the GPE.

A primary advantage of the imaginary time technique is the extreme simplicity of implementation, given the existence of a numerical solver for the real-time version of the equation. An important disadvantage is that the normalisation is controlled only indirectly via the chemical potential, and it is common to desire a fixed total number of atoms in a simulation instead. In the following we describe a method of achieving this by evolving μ along with the imaginary time evolution of the fields.

C.1 The usual imaginary time method

We start by investigating the reasons why the traditional imaginary time method works. In the case that the interaction strength is zero, we have the linear Schrödinger equation and the reasoning behind the method is straightforward: The coefficient of the lowest energy eigenstate evolves as $e^{-E_0\tau/\hbar}$ and therefore becomes large compared to coefficients $e^{-E_k\tau/\hbar}$ as $\tau \rightarrow \infty$, where E_k is the energy of the k th eigenstate. The wavefunction converges exponentially to the ground state as a result, provided that it has nonzero overlap with the initial state.

In the case that the Hamiltonian is nonlinear there is no simple formal solution to rely on, so the justification for the imaginary time method must be modified. The appropriate starting point is to realise that we seek to minimise the Hamiltonian under the constraint that the

number of atoms, N is fixed. That is, we wish to solve the constrained minimisation problem,

$$\begin{aligned} &\text{minimise } \mathcal{H}[\phi] \\ &\text{subject to } \mathcal{N}[\phi] = N \end{aligned} \tag{C.1}$$

where $\mathcal{N}[\phi] = \int dV |\phi|^2$ is the normalisation functional and N is the desired normalisation. Proceeding via Lagrange multipliers, we may define

$$\mathcal{K}[\phi] \equiv \mathcal{H}[\phi] - \mu (\mathcal{N}[\phi] - N), \tag{C.2}$$

and solve for the critical point of \mathcal{K} . It is important to note that the critical point of the Lagrange function is *not* a minimum with respect to the variables $\{\phi, \mu\}$; it is a saddle point. To find the critical point we must solve the two equations

$$\frac{\delta \mathcal{K}}{\delta \phi^*} = 0 \quad \text{and} \quad \frac{\partial \mathcal{K}}{\partial \mu} = 0. \tag{C.3}$$

The first of these is the condition that the right hand side of the GPE with an extra energy offset μ is zero, while the second recovers the number constraint:

$$(H_{\text{sp}} - \mu)\phi + U|\phi|^2\phi = 0 \quad \text{and} \quad \mathcal{N}[\phi] - N = 0. \tag{C.4}$$

In the conventional imaginary time method, we solve a restricted version of this problem by fixing μ and *minimising* \mathcal{K} . This works because the critical point of \mathcal{K} with respect to ϕ is a minimum for any fixed μ , even though it is not a minimum when the combined variables $\{\phi, \mu\}$ are considered. (To see this, note that \mathcal{K} is also a valid Hamiltonian with an energy offset μ .)

Minimisation of \mathcal{K} with fixed μ may be achieved by continuous-time steepest-descent optimisation and this corresponds to the imaginary time evolution. This may seem like a coincidence but the fact that imaginary time evolution minimises the energy is a generic property of Hamiltonian systems in complex phase space.

To see this, note that the GPE may be written

$$i\hbar \frac{\partial \phi}{\partial t} = \frac{\delta \mathcal{H}}{\delta \phi^*}. \tag{C.5}$$

On the other hand, the steepest descent direction for \mathcal{K} is $-\frac{\delta \mathcal{K}}{\delta \phi^*}$, so an evolution equation which takes us toward the ground state is

$$\hbar \frac{\partial \phi}{\partial \tau} = -\frac{\delta \mathcal{K}}{\delta \phi^*}. \tag{C.6}$$

The factor of \hbar on the left is included only to emphasise the similarities with Eq. (C.5). Taken

together, this shows why the imaginary time method works: it corresponds to a steepest-descent minimisation of the Hamiltonian, with chemical potential added as a Lagrange multiplier for the number constraint.

C.2 The imaginary time method with varying μ

The conventional imaginary time method as described above does not allow us to satisfy the normalisation constraint directly; to do this we must modify μ during the evolution. Surprisingly, the appropriate evolution law turns out to be the *uphill* evolution

$$\frac{\partial \mu}{\partial t} = + \frac{\partial \mathcal{K}}{\partial \mu}. \quad (\text{C.7})$$

Combined with Eq. (C.6) this takes us to the appropriate saddle point of \mathcal{K} representing the solution of Eqs. (C.3).

A discrete-time version of this saddle point finding algorithm is described in [82, §14.2 p.429] and the local convergence of the continuous version proven in Ref. [6, Theorem 1]. For convergence to hold, the Hessian of \mathcal{K} with respect to ϕ must be positive definite at the critical point. We note that this is the case for our system, because for every fixed μ the critical point is in fact a minimum with respect to variations of ϕ . While the convergence theorem holds only locally, we find that the Thomas-Fermi solution is a sufficiently accurate starting point for reliable convergence in practice.

We finish by giving some intuitive reasoning behind the saddle point algorithm. To simplify the exposition, we work with an analogous minimisation problem over \mathbb{R}^n :

$$\begin{aligned} &\text{minimise } f(\mathbf{x}) \\ &\text{subject to } g(\mathbf{x}) = 0 \end{aligned} \quad (\text{C.8})$$

where $f, g: \mathbb{R}^n \rightarrow \mathbb{R}$. The Lagrangian function for this minimisation is

$$\mathcal{L}(\mathbf{x}, \mu) = f(\mathbf{x}) - \mu g(\mathbf{x}), \quad (\text{C.9})$$

and we wish to solve for a critical point (\mathbf{x}_0, μ_0) of \mathcal{L} . In nearly all cases (\mathbf{x}_0, μ_0) is a saddle point because to first order near the critical point,

$$\mu g(\mathbf{x}) = \mu \left[g(\mathbf{x}_0) + \left. \frac{\partial g}{\partial \mathbf{x}} \right|_{\mathbf{x}_0} \cdot (\mathbf{x} - \mathbf{x}_0) + \dots \right] \quad (\text{C.10})$$

$$\approx \mu \nabla g(\mathbf{x}_0) \cdot (\mathbf{x} - \mathbf{x}_0), \quad (\text{C.11})$$

where the constant term disappears because the constraint g is satisfied at the critical point. This first-order approximation therefore amounts to a sum of saddle-like terms of the form $\mu(x_i - x_{0,i})$.

In the case that the principle directions of the saddle point correspond to the coordinate axes, for example

$$h_1(x, y) = x^2 - y^2, \quad (\text{C.12})$$

one can intuitively expect that a steepest descent in x and steepest ascent in y would converge to the saddle point. In the Lagrangian case, the principle directions do not line up with the coordinate axes, and instead we have a saddle point intuitively similar to the function

$$h_2(x, y) = x^2 - xy. \quad (\text{C.13})$$

Nevertheless, the evolution

$$\dot{\mathbf{x}} = -\frac{\partial \mathcal{L}}{\partial \mathbf{x}} \quad (\text{C.14})$$

$$\dot{\mu} = +\frac{\partial \mathcal{L}}{\partial \mu} \quad (\text{C.15})$$

still converges to the saddle point in many cases, provided we start sufficiently close. As noted above, this happens when the Hessian of \mathcal{L} with respect to \mathbf{x} is positive definite [6, Theorem 1]. This is the case when finding ground states of the GPE using the saddle point algorithm.

Additional work

The following two papers contain additional work completed during the course of the PhD, but which do not tie into the main theme of the thesis.

In the first paper we report the first discovery of a Bell inequality for observables with continuous spectra. In this paper the current author contributed the no-go proof on pages three and four, in collaboration with E. G. Cavalcanti: There are no local hidden variable inequalities possible if one considers only the first-moment correlations between continuous variables at different sites. The current author also helped with the final stages of writing the paper.

The second paper proposes a computationally efficient alternative to the traditional Uhlmann-Jozsa fidelity measure between two quantum states. The current author contributed section IV that discusses computational efficiency, including the optimised C implementations, and additional optimisation of the slower Matlab codes. The current author also helped with polishing of the manuscript as a whole.

Note for the arXiv version: Verbatim copies of the two papers, Refs. [18] and [84] were attached here in the version of this thesis submitted for examination. These have been omitted from the arXiv version due to technical constraints.